

GLOBAL PDF-BASED TEMPORAL NON-LOCAL MEANS FILTERING REVEALS INDIVIDUAL DIFFERENCES IN BRAIN CONNECTIVITY

Jian Li¹, Soyoung Choi^{2,1}, Anand A. Joshi¹, Jessica L. Wisnowski³, Richard M. Leahy¹

¹Signal and Image Processing Institute, University of Southern California, Los Angeles, 90089

²Neuroscience Graduate Program, University of Southern California, Los Angeles, 90089

³Radiology and Pediatrics, Division of Neonatology, Children’s Hospital Los Angeles | Keck School of Medicine, University of Southern California, 90089

ABSTRACT

Characterizing functional brain connectivity using resting fMRI is challenging due to the relatively small BOLD signal contrast and low SNR. Gaussian filtering tends to undermine the individual differences detected by analysis of BOLD signal by smoothing signals across boundaries of different functional areas. Temporal non-local means (tNLM) filtering denoises fMRI data while preserving spatial structures but the kernel and parameters for tNLM filter need to be chosen carefully in order to achieve optimal results. Global PDF-based tNLM filtering (GPDF) is a new, data-dependent optimized kernel function for tNLM filtering which enables us to perform *global* filtering with improved *noise reduction* effects *without blurring* adjacent functional regions.

Index Terms— non-local means, filtering, optimization, fMRI, connectivity

1. INTRODUCTION

Functional MRI (fMRI) is a powerful in-vivo neuroimaging tool that allows us to indirectly infer information about the neuronal activity of the brain by observing blood-oxygen level dependent (BOLD) signal fluctuations [1]. Temporal correlations in resting fMRI (rfMRI) BOLD signals across multiple spatially distinct brain regions are used to define functional brain networks [2]. However, BOLD signals inherently have low signal to noise ratio (SNR). Preprocessing of fMRI data often includes a spatial smoothing step to reduce noise. Isotropic 3D Gaussian filtering is the most commonly used approach to smooth volumetric rfMRI data [3], or equivalently, the Laplace-Beltrami (LB) operator is applied when the data is mapped onto a 2D representation of the cortical surface [4]. Both methods suffer from a critical common problem as they both spatially mix signals between adjacent functional regions, limiting our ability to accurately identify connectivity at the micro-to-meso scale in individual fMRI recordings.

Non-local means (NLM) filtering is an edge-preserving method originally designed for natural image denoising [5] and has been adapted to filter anatomical MRI [6], fMRI [7] and diffusion MRI [8] to preserve spatial structures in imaging data. Our laboratory recently developed a variant for filtering rfMRI data called temporal NLM (tNLM) that assigns non-local smoothing kernel weights based on temporal similarities between time series rather than spatial similarities [9]. We demonstrated tNLM filtering’s ability to

reduce noise by using (weighted) averages of only those time series that are similar, thus minimizing blurring across functional boundaries.

Here we identify two key challenges in using tNLM filtering as described in [9]. The first of these is that the exponential kernel function used in computing the weights is chosen heuristically. The exponent is an affine function of the sample correlation between the two time-series. As we show below, this function does not perform well in terms of optimizing the trade-off between using large weights when the correlations are large and smaller (or near zero) weights for low correlations. A second issue is that almost all NLM-based filtering methods, including tNLM, have been applied over a restricted neighborhood around the point to be filtered, partially because of the high computational cost if they are applied globally. However, since networks span the entire brain, global rather than local filtering has the potential for improved results when filtering using tNLM. It has been suggested previously that the brain has the structure of a small-world network [10] and therefore most “nodes” (or voxels) in the brain are not strongly correlated with each other. As a result, when filtering a particular node using data from the entire brain, the fraction of uncorrelated nodes is much larger than the portion of correlated nodes. This can result in an undue influence of the large number of uncorrelated nodes on the filtered signal if the filter weights applied to these nodes are not sufficiently suppressed. We address each of these issues in the method described below.

Here we propose Global PDF-based tNLM filtering (GPDF): a new kernel function for tNLM filtering of fMRI data based on the probability density function (PDF) of the correlation of the time series between pairs of voxels. This method enables us to perform *global* filtering with improved *noise reduction* effects while minimizing *blurring* of adjacent functional regions.

2. METHODS

2.1 NLM-based Filtering and tNLM

Let’s assume the fMRI data are represented on a 2D tessellation of the mid-cortical surface with V vertices and T time samples per vertex. Let $s(i, t)$ be the time series at vertex $i \in V$ and time $t \in T$. Let S_i be the set of vertices that are used to compute the filtered signal at vertex i . In the tNLM method, S_i contains vertex i and all its k -hop neighboring vertices, for some $k > 0$. Then tNLM filtering is defined as

$$s'(i, t) = \frac{1}{\sum_{j \in S_i} w(i, j)} \sum_{j \in S_i} s(j, t) w(i, j) \quad (1)$$

where the weight $w(i, j)$ is chosen to be a temporal similarity measure and defined as a function of the correlation [9]:

$$w(i, j) = f(r(i, j); h) \quad (2)$$

$$f_{\text{tNLM}}(r; h) = \exp\left(-\frac{2(1-r)}{h^2}\right) \quad (3)$$

where $r(i, j)$ is the Pearson correlation coefficient between vertices i and j and h is the parameter which controls the degree of filtering.

2.2 Global PDF-based tNLM Filtering (GPDF)

Our GPDF filtering differs from the original tNLM filtering in the following two ways: (i) the spatial range over which the filtered signal is computed: in GPDF the set $S_i = S, \forall i$, where S contains all vertices on the tessellated brain surface instead of just a local neighborhood; (ii) we use a different kernel function f in equation (2).

2.2.1 GPDF Kernel Formulation

Let the observed signal be $x_i = s_i + n_i$ at vertex i , a superposition of the true signal s_i and noise n_i . Assume that s_i and n_i are independent with $s_i \sim N(0, \sigma_s^2)$ and $n_i \sim N(0, \sigma_n^2)$. Also assume some non-zero correlation between s_i and s_j if i and j are within the same functional network (H_1) and zero correlation if they are in different networks (H_0). Then the correlation between two observed signals is in the form of:

$$\rho_{ij} = \frac{E[x_i x_j]}{\sigma_{x_i} \sigma_{x_j}} = \frac{E[s_i s_j]}{\sigma_s^2 + \sigma_n^2} = \begin{cases} 0 & , H_0: E[s_i s_j] = 0 \\ \frac{\sigma_s^2 c}{\sigma_s^2 + \sigma_n^2} & , H_1: \frac{E[s_i s_j]}{\sigma_s^2} = c \end{cases} \quad (4)$$

where $c \in [-1, 0] \cup (0, 1]$ represents some non-zero correlation and σ_{x_i} represents the standard deviation of x_i . To further help avoid numerical issues and improve the robustness of the algorithm described below, we formulate our hypothesis in a slightly relaxed form:

$$\rho_{ij} = \frac{E[x_i x_j]}{\sigma_{x_i} \sigma_{x_j}} \in \begin{cases} [-\delta, \delta] & , H_0 \\ [-1, -\delta) \cup (\delta, 1] & , H_1 \end{cases} \quad (5)$$

where δ is a small positive constant. The sample correlation distribution is given by the following [11]

$$P(r|\rho; T) = \frac{(T-2)\Gamma(T-1)(1-\rho^2)^{\frac{T-1}{2}}(1-r^2)^{\frac{T-4}{2}}}{\sqrt{2\pi}\Gamma\left(T-\frac{1}{2}\right)(1-\rho r)^{M-\frac{3}{2}}} \times {}_2F_1\left(\frac{1}{2}, \frac{1}{2}, \frac{2T-1}{2}, \frac{\rho r+1}{2}\right) \quad (6)$$

where T is the number of samples and ${}_2F_1(a, b; c; z)$ is the Gaussian hypergeometric function. The parameter T will be omitted in the following derivation as for a given fMRI dataset, T is a fixed scalar.

An example is shown in Fig. 1 where $\rho = 0.2$ under H_1 (blue curve) and $\rho = 0$ under H_0 (red curve). The histograms of the sample correlations are distributed about their means according to (6) due to the finite number of samples. This causes a significant overlap between the red and blue curves. There is therefore a range of nonzero correlation values over which it is difficult to distinguish H_1 from H_0 given an observed sample correlation r . But to perform well, tNLM

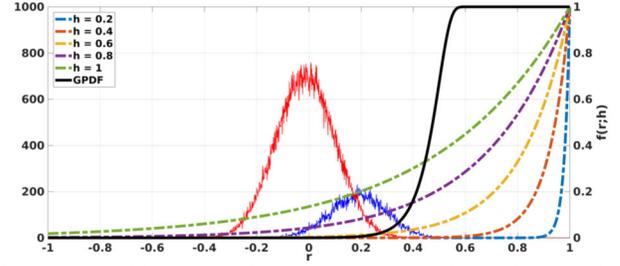


Figure 1: The histogram of the correlations under H_1 (blue) and H_0 (red) generated from simulated data overlaid with tNLM kernel functions for different parameter h (dotted) and GPDF kernel function (black solid).

should attach large weights only to those time series for which H_1 is true.

In Fig. 1 we show the shape of the original tNLM kernel defined in (3) as a function of h . The figure shows that the kernel performs a poor job in differentiating H_1 from H_0 , in the sense that applying significant weights for H_1 also results in weights significantly greater than zero for H_0 . The black curve shows an alternative kernel that, visually at least, does a better job of giving significantly larger weights to H_1 while minimizing those for H_0 . We now describe how we select this kernel and then evaluate its performance.

Bayes theorem tells us the posterior probability of ρ given r is

$$P(\rho|r) = \frac{P(r|\rho)P(\rho)}{\int P(r|\rho)P(\rho)d\rho} \quad (7)$$

To better differentiate H_1 from H_0 , we take the ratio between the integrated posterior probability under H_1 and the counterpart under H_0 , forming the Bayes factor [12]

$$R(r) = \frac{\int_{\rho \in H_1} P(r|\rho)P(\rho)d\rho}{\int_{\rho \in H_0} P(r|\rho)P(\rho)d\rho} \quad (8)$$

where $R(r) \in [0, \infty)$. The larger $R(r)$, the more likely ρ belongs to H_1 given that sample correlation r .

We then reformulate our kernel function f to be

$$f_{\text{GPDF}}(r; h) = 1 - \exp\left(-\frac{R(r)}{h^2}\right) \quad (9)$$

where, similar to the tNLM kernel in equation (3), h is a parameter that controls the degree of smoothing. Replacing the sample correlation in (3) with the Bayes factor in (8) introduces the strong nonlinearity visible in the black curve in Fig. 1. This nonlinearity accounts for the fact that the posterior probability of H_1 vs H_0 can change rapidly as a function of r , as reflected in the Bayes factors.

2.2.2 Automated Parameter Selection

In addition to using a different kernel, we also propose an automated method for selecting the parameter h . To do this we maximize the expected value of the weighting function $f_{\text{GPDF}}(w; h)$ under H_1 while controlling the mean value with respect to H_0 . Specifically,

$$\hat{h} = \arg\max_h E_{H_1}\{f(r; h)\}, \text{ s.t. } E_{H_0}\{f(r; h)\} \leq \alpha \quad (10)$$

where α is the expected weight under H_0 , analogous to the false positive rate in detection theory. Although α is another parameter we need to tune manually, it is much more meaningful and robust than h , because choosing the same α

will generally yield similar filtering result while the internal parameter h can have very different impact for different datasets, as a function of the noise level, range of correlation values and size of the image being filtered. We recommend that α be set conservatively, e.g. 10^{-3} or smaller, due to the dominant volume of uncorrelated vertices in an fMRI dataset.

2.2.3 Estimation of the Population Correlation Distribution

In order to construct the kernel function in equation (9) we need to know the Bayes factor $R(r)$, which requires the conditional distribution $P(r|\rho)$ and the population correlation distribution $P(\rho)$. The sample correlation density $P(r|\rho)$ has an analytical solution given in equation (6). Therefore, we need only to estimate $P(\rho)$. Let $P(r)$ be the empirical sample correlation distribution obtained from the fMRI data. Let $P'(r) \in \mathbb{R}^M$, $P'(r|\rho) \in \mathbb{R}^{M \times N}$ and $P'(\rho) \in \mathbb{R}^N$ be the discretized version of the corresponding variables in the continuous space, respectively. Then $P'(\rho)$ can be estimated using a linear regression with non-negative constraints.

$$\hat{P}'(\rho) = \underset{P'(\rho)}{\operatorname{argmin}} \|P'(r) - \sum_{\rho} P'(r|\rho)P'(\rho)\|_{l_2}^2, \text{ s.t. } P'(\rho) \geq 0 \quad (11)$$

This optimization is a well-posed problem as long as $M \geq N$, i.e. the discretization step for ρ is smaller than that for r , which can be achieved easily. Also, this problem can be solved efficiently using the non-negative least square method.

2.2.4 GPDF Filtering Algorithm

We summarize our GPDF filtering algorithm as follows:

1. Given fMRI data $X \in \mathbb{R}^{V \times T}$, calculate the correlation matrix $A = XX^T \in \mathbb{R}^{V \times V}$.
2. Estimate $P'(r)$ from the histogram of the elements of A .
3. Estimate the priors by solving equation (11)
4. Optimize the parameter h by solving equation (10)
5. Construct the kernel using equation (9)
6. Finally filter the signal using equation (1)

3. EXPERIMENTS AND RESULTS

3.1 Simulation

We simulated the tessellation of the brain surface with 2D blocks of size $V \times V$ ($V = 32$) representing left and right hemispheres. Each point in each block represents a vertex on the brain surface and has a label, indicating which network it belongs to. Fig. 2 (a) shows the ground truth label blocks where each color represents a distinct label. The top and bottom rows have identical labels to simulate connections between the right and left hemispheres (in total $K = 16$ unique labels). For each label, we generated a random time series (white noise) of length $T = 200$ where points within the same labels were given identical time series (perfectly correlated) in the absence of noise. Points with different labels were given zero correlation indicating that they belong to different networks. We then added Gaussian white noise with $SNR = 0.4$ to the entire dataset.

To investigate the effects of different filtering methods, we applied filtering to the simulated data then parcellated the data into K labels using normalized cuts (Ncuts) [13]. A stable matching algorithm [14] was applied to match labels

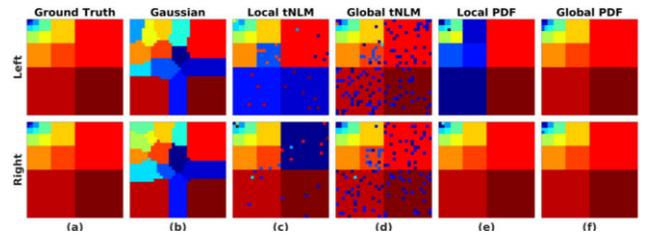


Figure 2: Parcellation result of simulated data represented as a $V \times V$ matrix for each method and each hemisphere. Columns from (a) to (f) are indicated by their titles along upper row. The rows represent the two hemispheres.

between different results for easy comparison. Figure 2 displays the parcellation results for: (b) Gaussian filtering with full-width-half-maximum (FWHM) approximately 8 points; (c), (d) tNLM filtering with optimized h parameter [15]; (e), (f) PDF filtering. To demonstrate the difference between local filtering and global filtering, we applied tNLM and PDF both locally ((c) and (e)) and globally ((d) and (f)). Local filtering processed left and right hemispheres separately while global filtering processed them jointly.

Gaussian spatial filtering generated labels along the boundaries between true labels not seen in the ground truth. This is most likely due to blurring of uncorrelated, neighboring vertices. In contrast, both tNLM and PDF filtering methods preserved the blocky structures. However, PDF yielded much cleaner results than tNLM because tNLM has a larger contribution from the uncorrelated vertices at each filtered point in order to assign higher weights to the correlated points, as discussed above. Note that for both PDF and tNLM the parameter h had been optimized, in the latter case using [15], to achieve the best trade-off. Finally, for both tNLM and PDF, local filtering resulted in labels that were mismatched between the left and right hemispheres. The myopic perspective of local filtering failed to detect the distal inter-hemispheric connections.

We also ran this simulation for 100 Monte Carlo trials and calculated the Adjusted Rand Index (ARI) [16] between each parcellation result and the ground truth as a filtering performance measure. Results showed that the medians of the ARIs were 0.547, 0.701, 0.760, 0.750, 0.969 respectively in correspondence to each filtering method in Fig. 2 (b) – (f), respectively, indicating that GPDF outperformed other filtering methods by a significant margin.

3.2 Application to rfMRI Dataset

3.2.1 Dataset and Filtering

40 subjects with minimally preprocessed rfMRI datasets (2 sessions, 2 phase encodings; 160 sessions total) were obtained from Human Connectome Project (HCP) [17]. The data were acquired with $TR=720\text{ms}$ with resolution $2 \times 2 \times 2$ mm and had been carefully preprocessed using the pipeline described in [18]. Then the data were co-registered onto a common atlas and downsampled onto a 32K-vertex cortical surface. We further downsampled each data to 11K vertices for computational tractability.

We then filtered each dataset using LB with $\sigma = 2\text{mm}$ and GPDF with $\alpha = 10^{-3}$. We did not include tNLM filtering results here because tNLM had been extensively

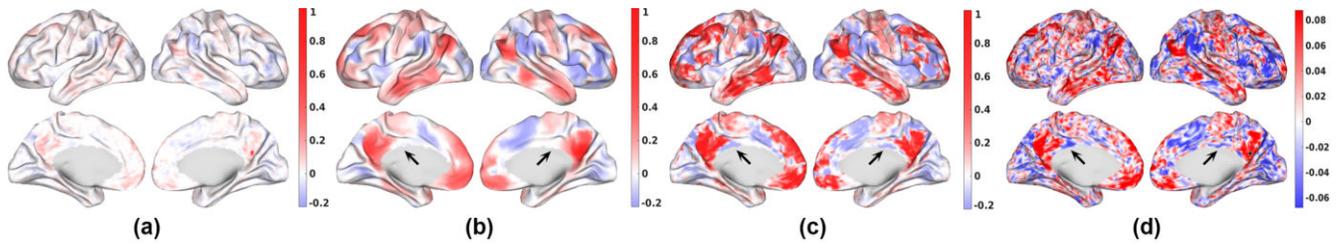


Figure 3: Seeded correlation map for a single subject for (a) unfiltered data; (b) LB filtered data; (c) GPDF filtered data; (d) unfiltered data re-plotted in its own scale. Seed point was selected in the caudal pre-cuneus area shown as a black dot in the bottom right sub-figure of (d). Positively correlated regions are shown in red, uncorrelated regions in white and negatively correlated regions in blue.

studied and compared with LB on real datasets in [9] and [15].

3.2.2 Seeded Correlation Map

To *qualitatively* evaluate the effects of filtering, we used a seed point in the pre-cuneus which is part of the default mode network (DMN) (Fig. 3 (d)) and calculated its correlation with all other vertices of the brain, forming a correlation map. Fig. 3 shows seed-point correlation maps for a single subject for (a) unfiltered data; (b) LB filtered data and (c) GPDF filtered data in a common scale ranging from -0.2 to 1 .

While low correlations were observed across the brain in unfiltered data due to rfMRI’s inherent low SNR, positive correlations with the regions of the DMN were observed. Figure 3 (d) exaggerates the color scale of unfiltered data for easy visualization of these spatial structures. LB and GPDF, in contrast, yielded higher correlations due to their ability to reduce noise and amplify signal. However, GPDF exhibited a wider range and stronger correlation values than LB.

Additionally, GPDF appears better able to preserve spatial structures between adjacent ROIs (functional regions) with opposite correlation values to the seed point. The boundary between two adjacent functional regions are indicated by the arrows in Fig. 3. This boundary is observed in both unfiltered data and GPDF but not in LB. These observations are indicative of LB’s tendency to spatially blur the boundaries between adjacent ROIs.

LB showed strong connections to the local points surrounding the seed point while connections to distal areas, especially inter-hemispherical connections, were strongly attenuated due to the localness of the filtering. This attenuation did not occur in GPDF as strong correlations are preserved across distal and inter-hemispheric regions of the DMN. GPDF therefore appears to help reveal stronger intra-network connectivity than the LB filtering method.

3.2.3 Unfiltered Correlation Matrix and Modularity

To further *quantitatively* evaluate the filtering performance, for each dataset we took the unfiltered data and computed the vertex-pairwise full correlation matrix, $A \in \mathbb{R}^{V \times V}$ and binarized it with threshold th to form a binary adjacency matrix A' . We also applied the Ncuts algorithm to parcellate the brain into K networks using each of the following: the unfiltered data, the LB-filtered data and the PDF-filtered data. We then calculated the modularity [19] for the adjacency matrix A' using each of the three K network partitions (unfiltered, LB, PDF) as a function of threshold th . Using the same unfiltered data adjacency A' establishes an unbiased comparison of the three partitions. The resulting modularity

measure indicates how well each filtering method grouped the data into functionally homogeneous regions with respect to the original (unfiltered) data. In essence, we assume that a better filtering method will give us a better clustering of the nodes under a given parcellation algorithm, in the sense that nodes that have the same labels (are within the same network) tend to have higher and consistent correlation with each other than with nodes in other networks.

The analyses above were performed on each dataset independently. Fig. 4 shows the median modularity across 160 sessions (40 subjects \times 4 sessions) as a function of the threshold th . The GPDF filtering method outperformed LB and the unfiltered case by a large margin regardless of the threshold settings and the number of parcels, indicating that GPDF is producing parcellations that show stronger within network similarity in the raw (unfiltered) data than either unfiltered data or LB filtering. It also worths noting that LB filtering actually performs worse than the unfiltered case when performing individual parcellations, suggesting that LB may not optimally preserve differences between individuals based on a single fMRI recording.

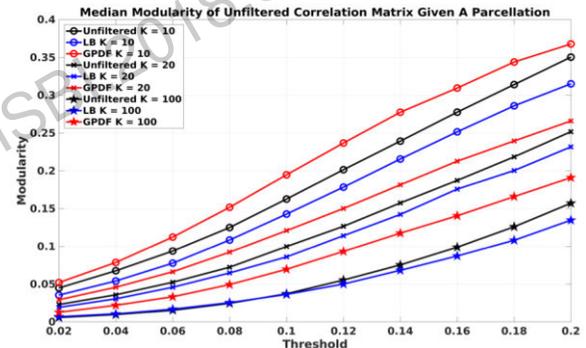


Figure 4: Modularity (y-axis) as a function of threshold value (x-axis) for unfiltered parcellation result (black), LB filtered (blue) and GPDF filtered (red). Different number of parcels K are shown with different markers: circle for $K = 10$, x-mark for $K = 20$ and pentagram for $K = 100$.

4. CONCLUSION

In this paper, we developed a novel kernel function for global tNLM filtering. We have demonstrated qualitatively and quantitatively that this method can perform better denoising than standard linear filtering method when the filtering is performed for the purposes of network identification. The approach may be particularly useful when inferring connectivity patterns from individual fMRI recordings. Extensions of this method to explore dynamic functional connectivity is a promising future direction.

REFERENCES

- [1] B. Biswal, Y. FZ, H. VM, and H. JS, “Functional connectivity in the motor cortex of resting human brain using,” *Magn Reson Med*, vol. 34, no. 9, pp. 537–541, 1995.
- [2] S. M. Smith *et al.*, “Correspondence of the brain’s functional architecture during activation and rest,” *Proc. Natl. Acad. Sci. U. S. A.*, vol. 106, no. 31, pp. 13040–5, 2009.
- [3] S. M. Smith *et al.*, “NeuroImage Resting-state fMRI in the Human Connectome Project,” *Neuroimage*, vol. 80, pp. 144–168, 2013.
- [4] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis, “On the Laplace-Beltrami operator and brain surface flattening,” *IEEE Trans. Med. Imaging*, vol. 18, no. 8, pp. 700–11, 1999.
- [5] A. Buades, B. Coll, and J.-M. Morel, “A non-local algorithm for image denoising,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2005, vol. 2, pp. 60–65.
- [6] J. V Manjon, J. Carbonell-Caballero, J. J. Lull, G. Garcia-Marti, L. Marti-Bonmati, and M. Robles, “MRI denoising using non-local means,” *Med Image Anal*, vol. 12, no. 4, pp. 514–523, 2008.
- [7] 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,” *Front. Hum. Neurosci.*, vol. 8, no. September, p. 715, 2014.
- [8] N. Wiest-Daesslé, S. Prima, P. Coupé, S. P. Morrissey, and C. Barillot, “Rician noise removal by non-local means filtering for low signal-to-noise ratio MRI: applications to DT-MRI,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2008, pp. 171–179.
- [9] C. Bhushan *et al.*, “Temporal non-local means filtering reveals real-time whole-brain cortical interactions in resting fMRI,” *PLoS One*, vol. 11, no. 7, pp. 1–22, 2016.
- [10] E. Bullmore and O. Sporns, “Complex brain networks: graph theoretical analysis of structural and functional systems,” *Nat. Rev. Neurosci.*, vol. 10, no. 4, pp. 312–312, 2009.
- [11] R. a. Fisher and R. a. Fisher, “Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population,” *Biometrika*, vol. 10, no. 4, pp. 507–521, 1915.
- [12] R. E. Kass and A. E. Raftery, “Bayes factors,” *Journal of the American Statistical Association*, vol. 90, no. 430, pp. 773–795, 1995.
- [13] J. Shi and J. Malik, “Normalized cuts and image segmentation,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 8, pp. 888–905, 2000.
- [14] D. Gale and L. S. Shapley, “College Admissions and the Stability of Marriage,” *Am. Math. Mon.*, vol. 69, no. 1, pp. 9–15, 1962.
- [15] J. Li and R. M. Leahy, “Parameter selection for optimized non-local means filtering of task fMRI,” *Proc. - Int. Symp. Biomed. Imaging*, pp. 476–480, 2017.
- [16] W. M. Rand, “Objective criteria for the evaluation of clustering methods,” *J. Am. Stat. Assoc.*, vol. 66, no. 336, pp. 846–850, 1971.
- [17] H. WU-Minn, “500 Subjects+ MEG2 Data Release: Reference Manual,” 2014.
- [18] M. F. Glasser *et al.*, “The minimal preprocessing pipelines for the Human Connectome Project,” *Neuroimage*, vol. 80, pp. 105–124, 2013.
- [19] M. E. J. Newman, “Modularity and community structure in networks,” *Proc. Natl. Acad. Sci.*, vol. 103, no. 23, pp. 8577–8582, 2006.