

DENOISING OF fMRI VOLUMES USING LOCAL LOW RANK METHODS

Pierre-Antoine Comby^{*†}

Zaineb Amor^{*}

Alexandre Vignaud^{*}

Philippe Ciuciu^{*†}

^{*}CEA, Joliot, NeuroSpin, Université Paris-Saclay, F-91191 Gif-sur-Yvette, France

[†]Inria, MIND, Université Paris-Saclay, F-91120 Palaiseau, France

ABSTRACT

Functional Magnetic Resonance Imaging (fMRI) data provides deep insight on brain activity, but high-resolution (e.g. 1 mm isotropic) fMRI suffers from low signal-to-noise ratio (SNR). Recently several denoising methods have been proposed to boost the SNR of high-resolution fMRI images. In this work we carry out a prospective benchmark of local low-rank denoising methods by quantifying their performances on a retinotopy experimental paradigm that was submitted to six healthy volunteers during a 7 Tesla, 1 mm isotropic fMRI acquisition protocol. Out of the five tested approaches (NORDIC, MP-PCA, Hybrid-PCA, Optimal Threshold (OT) and Hybrid-OT), the OT approach outperforms its competitors in terms of statistical sensitivity and specificity both at the subject- and group-level. Additionally, by testing these denoising methods in different configurations of preprocessing pipelines, we demonstrate first that on average it is beneficial to denoise fMRI images prior to performing realignment (i.e. motion correction), and second that the OT approach performs better when applied to complex-valued fMRI images instead of magnitude-only ones. We also provide an open source implementation to promote a broader use of denoising methods in fMRI and enable reliable statistical data analysis at high spatial resolution.

Index Terms— functional MRI, patch denoising, singular value thresholding

1. INTRODUCTION

Functional Magnetic Resonance Imaging (fMRI) uses the blood oxygenation level dependent (BOLD) contrast as a proxy to study local neuronal activity [1], and remains one of the preferred imaging techniques to study brain function in healthy and pathological condition [2]. In this extend fMRI seeks to maximize the BOLD contrast between brain states (activation vs baseline), hence, the available temporal SNR, resulting from a compromise between a targeted spatio-temporal resolution and physical acquisition constraints. Using current state-of-the-art accelerated imaging techniques such as Parallel Imaging [3] and Compressed Sensing (CS) [4] as well as ultra high magnetic field system (7 Tesla, 7 T) [5], it is now feasible to collect fMRI data at submillimetric resolution with reasonable volumetric repetition time ($TR \simeq 1.2s$).

However with such settings, the fMRI data is heavily spoiled by surrounding noise sources: Firstly the thermal noise emanating from the acquisition process contaminates the fMRI signal as a complex independent Gaussian source with spatially varying variance [6]. Secondly the physiological noise, i.e. patient-induced signal variations that are acquired along with the signal of interest but do not reflect significant brain activity (breathing, heart beat, head motion, etc.) [7] (3D-EPI sequences are more sensible to physiological noise than the classical 2D EPI scheme).

Overall, these noise sources limit the reliability of fMRI data and constrain neuroscientists to collect multiple runs in each individual to boost the statistical sensitivity.

To increase the statistical significance of an fMRI experiment, the sequence of fMRI images is usually pre-processed using state-of-the-art toolboxes such as FSL [3] and SPM [8]. Recently fMRIprep [9], leveraging the Python based Nipype package [10] has been proposed as a unified and standardized workflow.

Concomitantly denoising methods have been introduced for fMRI (or other dynamical imaging such as DWI) [11]–[19] and all assume gaussian distributed noise. Yet, their integration in the processing pipeline has hardly been considered, at most the denoising step is performed prior to any processing, or simply tested on synthetic data. Moreover, only the recent contribution of NORDIC [14], which rooted the interest for this work, considered complex-valued data, which preserve the gaussian distributed noise hypotheses, instead of using magnitude only images, where the noise is rician [20].

This paper aims at benchmarking denoising methods for fMRI, by quantifying the gain they provided in the downstream statistical analysis. In particular we evaluated 5 methods based on the local low-rank property of fMRI images (presented in Section 2) and testing how they can be optimally integrated with data preprocessing. To this end, we acquired 3D-EPI scans in six healthy volunteers undergoing a retinotopic mapping experiment (cf. Section 3), and applied a denoising step either on complex-valued or magnitude-only fMRI images. In Section 4 the results of standard statistical analysis are summarized over the six individuals. Conclusions are drawn in Section 5.

2. DENOISING METHODS USING PATCH BASED PCA

2.1. From global to local low rank assumption.

Let $\mathbf{Y}$ be the *complex-valued* fMRI scans obtained after reconstruction and combination of multi-coil data: $\mathbf{Y} = [\mathbf{y}_1 \dots \mathbf{y}_t \dots \mathbf{y}_{N_t}] \in \mathbb{C}^{N_x N_y N_z \times N_t}$, where $N_x N_y N_z$ is the number of voxels and N_t the number of scans. Along with the fMRI data, we can also retrieve a noise variance map $\Sigma \in \mathbb{R}_+^{N_x N_y N_z}$.

fMRI scans are inherently low rank as we observe small variations of the BOLD signal over a constant background. Hence, temporal frames are highly correlated. Furthermore, the periodic behavior of the physiological activities and the presence of repeated stimuli strengthen this effect. In such settings, the data at hand can be considered as a noisy observation of a low rank matrix $\mathbf{Y}_{ij} = \mathbf{X}_{ij} + \mathbf{N}_{ij}$ where the noise components are independent and Gaussian $\mathbf{N}_{ij} \sim \mathcal{N}(0, \Sigma^2)$. The low rank matrix is the solution of the problem [21]:

$$\hat{\mathbf{X}} = \arg \min_{\mathbf{X}} \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda \|\mathbf{X}\|_* \quad (1)$$

where $\|\mathbf{X}\|_* = \sum_{k=0}^n \sigma_k(\mathbf{X})$ denotes the nuclear norm of $\mathbf{X}$.

The analytical solution of (1) consists in applying a soft thresholding operator to the singular values of $\mathbf{Y}$. However, this global low rank modeling suffers in practice from the high dimensionality of the problem:

- (i) In the case of 3D fMRI acquisitions, we typically have $N_x N_y N_z \sim 10^5 N_t$ and the limited amount of degrees of freedom on the SVD will reduce the effectiveness of a rank constraint.
- (ii) The spatial noise level is heterogeneous in the context accelerated imaging due to multi-coil interactions [22].

To alleviate those challenges, problem (1) can be solved locally on small 3D+time patches extracted from the whole 4D fMRI sequence. The patch size and their overlap become hyperparameters for the problem.

2.2. Local low rank formalism

The patch extraction operator $\mathcal{P}_u(\mathbf{Y})$ basically extracts a $K = k_x \times k_y \times k_z$ dimensional patch centered in voxel $\mathbf{u}$ across all fMRI scans in $\mathbf{Y}$, yielding a $K \times N_t$ so-called Casorati matrix (see Fig. 1). $\mathbf{Y}_{[u]} = \mathcal{P}_u(\mathbf{Y})$ is a low rank matrix containing background (e.g. a T_2^* -w version of brain anatomy) information polluted by a zero-mean Gaussian noise of local variance $\sigma_{[u]}^2$: $\mathbf{Y}_{[u]} = \mathbf{X}_{[u]} + \mathbf{N}(\sigma_{[u]}^2)$.

Retrieving a low rank approximation of the patch typically consists in applying a singular value decomposition (SVD) of $\mathbf{Y}_{[u]} = \mathbf{U}\mathbf{S}\mathbf{V}^T$, and then thresholding it to retain the largest $n_{[u]}$ singular values related to the signal of interest and yield the low rank approximation $\hat{\mathbf{X}}_{[u]} = \mathbf{U}\mathbf{S}^* \mathbf{V}^T$, cf. Fig. 1. Once the noise has been discarded, the patches are recombined with a weighting mechanism in case of overlapping patches as originally proposed in [11]. Each voxel at position i in the final denoised fMRI sequence, is computed from the P patches containing it as follows:

$$\hat{\mathbf{X}}(i) = \frac{\sum_{j=1}^P w_j \hat{\mathbf{X}}_{[u_j]}(i)}{\sum_{j=1}^P w_j}, \quad w_j = \frac{1}{1 + n_{[u_j]}} \quad (2)$$




Fig. 1: General procedure for LLR denoising. Without loss of generality a 2D case is presented. The sequential data is processed into patches, that are (1.) extracted, (2.) SVD decomposed and thresholded, and then recombined (3.) using a weighted average in case of overlap.

2.3. Comparison of local low-rank methods

Several local low rank (LLR) denoising methods [11]–[19] have been proposed in the literature. We selected four out of them (cf.

Table 1) given their impact in the recent literature and their diversity to address the thresholding problem. When available, we directly used the original authors’ implementation. Each method proposed a thresholding function η over the singular values $\lambda_1 \dots \lambda_{N_t}$ of the patch $\mathbf{Y}_{[u]}$. For instance, the optimal threshold (OT) with respect to the Frobenius norm [17] is defined as follows:

$$\eta(\lambda) = \frac{N_t \sigma^2}{\lambda} \sqrt{\left(\frac{\lambda^2}{N_t \sigma^2} - \beta - 1 \right)^2 - 4\beta} \mathbf{1}_{\frac{\lambda}{\sqrt{N_t} \sigma} \geq 1 + \sqrt{\beta}} \quad (3)$$

with $\beta = N_t/K$, if a noise map Σ is available, then $\hat{\sigma}$ is the average of $\Sigma_{[u]}$ (referred as Hybrid-OT hereafter), else we resort to the robust estimator $\hat{\sigma} = \text{med} \lambda / \sqrt{N_t \mu_\beta}$, and μ_β is the median of the Marcenko-Pastur’s law [23], which describes the asymptotic distribution of the singular values of a normally distributed random matrix of aspect ratio β . This distribution is thus at the heart of the selected methods. Noticeably the adaptive thresholding technique proposed in [18] was not considered in the benchmark as its computational time was two order of magnitude larger than its LLR competitors. We also discarded Non-Local-Mean (NLM) based methods [24], [25], as they do not fit in the LLR framework. In Table 1, Hybrid-PCA and Hybrid-OT are variations of MP-PCA and OT approaches, respectively, in which a noise map (i.e. a covariance matrix Σ) estimate was provided.

Table 1: LLR methods under study. Defining the threshold for each patch in Hybrid-PCA requires an external noise variance map, not required by MP-PCA. NORDIC uses a global threshold, and normalizes the noise variance of each patch to apply it. Optimal-Threshold is described in more detail in text, cf Eq. (3).

Name	Thresholding $\eta(\lambda)$	Extra Data used
Nordic [14]	$\max(0, \lambda - \theta_{\text{global}})$	None
MP-PCA [12]	$\max(0, \lambda - \theta_{[u]})$	None
Hybrid-PCA [13]	$\max(0, \lambda - \theta_{[u]})$	Σ
OT [17]	Eq. (3)	None
Hybrid-OT [17]	Eq. (3)	Σ

For ease of comparison, we settle a patch size of $11 \times 11 \times 11$ voxels with an overlap of 5 voxels in each direction. Adding overlap reduces boundary effects, at the cost of larger processing time. In order to mitigate this effect, we only computed patches that shared at least 10% of their voxels with the brain mask, halving roughly the computational time.

3. MATERIAL AND METHODS

3.1. Acquisition

Six healthy volunteers were scanned on a Siemens Magnetom 7 T (Siemens-Healthineers, Erlangen, Germany) and a 1Tx-32Rx head coil (Nova Medical, Wilmington, CO, USA) using a 3D EPI sequence (1 mm – iso TE=20 ms, volumetric TR=2.4 s, 120 reps). Task-based fMRI data was collected along two consecutive runs during a retinotopic mapping paradigm, implemented in ¹, with a rotating wedge (clockwise and anti-clockwise) with a period of 32 s [26]. This paradigm promotes well localized BOLD signal in

¹https://github.com/hbp-brain-charting/public_protocols

the visual areas, however the spatial and temporal resolution considered were challenging in regard to the temporal SNR. The LLR denoising methods should thus bring significant benefit to the downstream task of fMRI data analysis in order to detect evoked brain activity and reconstruct the retinotopic phase maps.

3.2. Preprocessing

To quantify the potential benefit of each denoising method we tested 4 distinct preprocessing pipelines that consist in applying the denoising step either on the complex-valued (CD) or magnitude-only (MD) fMRI images and then in interchanging its position in the workflow (i.e. before or after image realignment (R) for motion correction). In case of complex denoising after realignment (i.e. R+CD scenario) the motion correction estimated from the magnitude images was applied to the real and imaginary parts before denoising. In the opposite scenario (CD+R), we computed the magnitude of fMRI images prior to performing motion correction.

Additionally, a denoise-only scenario was tested, but it underperformed compared to the standard workflow that embeds realignment. In every case the fMRI images are corrected for B_0 field inhomogeneities afterwards. To favor reproducible research, we provide open sourced implementations of the LLR denoising methods² and of the preprocessing workflow³.

3.3. Statistical Analysis

After testing the different configurations of the preprocessing pipeline, the fMRI images data were analyzed using the Nilearn package [27]. First, a general linear model (GLM) that embodies the two runs associated with the retinotopic experimental paradigm (clockwise and counter-clockwise) was built up. It includes 2 paradigm-related regressors (parametric, continuous and sinusoidal), 6 rigid motion regressors, a drift regressor and the baseline. The global effect of interest was determined using a statistical F -test over the two sinusoidal regressors ($H_0 : \alpha_{1,i}^2 + \alpha_{2,i}^2 = 0$), providing a z-score in each voxel i . To detect evoked brain activity we then thresholded the z-score map at $p < 0.05$, to correct for multiple comparisons using the false discovery rate (FDR).

4. RESULTS

4.1. Single subject analysis

In Fig. 2, we compare the z-score distributions related to the global effect of interest associated with the original and denoised fMRI images for a single subject. LLR denoising methods (color-coded) reach higher z-scores compared to the noisy baseline (in gray) whatever the type of fMRI images processed (magnitude-only or complex-valued). We selected this subject as that showing the best improvement between the baseline and the denoised versions.

It is worth noting that the type of input fMRI images slightly impacts the tail of the z-score histograms (see Fig. 2) as the noise statistics becomes Rician for magnitude-only images.

Moreover, the order of preprocessing steps (R: realign, D: denoising) matters when applied to complex-valued fMRI images: In that context, denoising must precede realignment as this combination provides the largest gain in statistical sensitivity. In contrast, for this subject the best order of preprocessing steps applied to magnitude-only fMRI images is less clear. Hence, to obtain more

reliable results, we performed the same analysis on the 6 participants. However, prior to analyzing the results at the group level (cf. Section 4.2), we pay attention to the spatial effect of denoising on the statistical z-score maps for this subject.




Fig. 2: Non-zero Z-score probability distribution (ν) for a single participant (subject-3). Top vs Bottom: Analysis of magnitude-only vs complex-valued fMRI images. The distribution has been smoothed with a Gaussian kernel for visualization purposes.

In Fig. 3 we show the denoised z-score maps yielded by the 20 tested pipelines as well as the noisy baseline. Since we only report surviving voxels to thresholding, we bring evidence that NORDIC, MP-PCA and Hybrid-PCA produced more false positives as they retrieve activations in the white matter on top of the expected ones in the visual cortex. Indeed the retinotopic paradigm is known to elicit evoked activity only in the occipital cortex.

Hence, OT-based approaches are more specific as their activations remain located in the gray matter. At the same time, the OT approach yields the highest z-scores with the largest spatial extent in the occipital region (cf Fig. 4), notably when the denoising is performed prior to realignment (scenario OT/MD+R). In this subject, we noticed that the behavior of Hybrid-PCA and MP-PCA is similar, and those approaches underperform compared to NORDIC and OT-based methods.

In Fig. 4 we actually zoomed in the occipital cortex and show first that all LLR denoising methods increase the spatial extent and significance of z-scores, hence confirming the whole brain results depicted in Fig. 2. Further, we demonstrate that the OT approach provides the best sensitivity/specificity trade-off even though we cannot access to the ground truth as we are not in a simulated framework.

4.2. “Group-level” analysis

In Table 2 we report the average gain across subjects for all tested denoising pipelines as a multiplicative factor MF_s obtained by dividing for each subject s the number of activating voxels associated with each LLR method over the count of activating voxels for the baseline (no denoising). Then we averaged these factors across subjects to get $\overline{MF} = 1/6 \sum_{s=1}^6 MF_s$. We computed this multiplicative factor first over the whole brain but only for the voxels

²<https://github.com/paquiteau/patch-denoising>

³<https://github.com/paquiteau/retino-pipeline>




Fig. 3: Z-score activation maps (axial slices) for the same participant as in Fig. 2. Only non-zero values are displayed.

surviving to FDR-corrected thresholding and then within a region of interest (ROI) located in the occipital cortex.

When comparing columns in Table 2, we can see that in almost every case, performing the denoising step prior to realignment is the most beneficial whether this treatment is applied to magnitude-only (MD+R) or to complex-valued (CD+R) fMRI images. This result is particularly significant when focusing on the occipital ROI.

Even though the realignment step aims at producing lower-rank fMRI images notably by removing part of the physiological noise, this step also resamples the images, thus potentially affecting the noise behavior and making it no longer additive. This could explain why the denoising step is less efficient when performed after motion correction, i.e. in R+MD vs MD+R or in R+CD vs CD+R.

In contrast to the claims done in [14], [15], our benchmark summarized in Table 2 shows that NORDIC on average possesses the smallest gain in activation (smallest $\overline{MF}$ value) compared to the baseline and its competitors. This is notably true when the denois-




Fig. 4: Zoom in the occipital cortex of z-score activation maps (axial slices) for the same participant as in Fig. 2, thresholded at $p < 0.05$ using FDR correction for multiple comparisons.

Table 2: Average gain (multiplicative factor) $\overline{MF}$ of activated voxels using denoising methods compared to baseline. Over 6 subjects and test cases, the highest gain is achieved by OT (**x8.03**) in the R+CD configuration. Whole brain vs ROI-based $\overline{MF}$ are reported as top/bottom numbers in each cell of the table.

Denoiser	NORDIC	MP-PCA	Hybrid-PCA	OT	Hybrid-OT
R+MD	$\times 3.52$	$\times 6.02$	$\times \mathbf{6.09}$	$\times 3.73$	$\times 0.91$
	$\times 3.64$	$\times 4.92$	$\times \mathbf{4.93}$	$\times 4.33$	$\times 1.00$
R+CD	$\times 0.57$	$\times 2.98$	$\times 3.27$	$\times \mathbf{8.03}$	$\times 5.29$
	$\times 0.55$	$\times 2.70$	$\times 2.91$	$\times \mathbf{5.04}$	$\times 4.01$
MD+R	$\times 3.36$	$\times 6.32$	$\times \mathbf{7.55}$	$\times 3.10$	$\times 1.22$
	$\times 3.48$	$\times 4.77$	$\times \mathbf{6.10}$	$\times 3.19$	$\times 4.90$
CD+R	$\times 2.57$	$\times 5.53$	$\times 4.97$	$\times \mathbf{7.91}$	$\times 5.45$
	$\times 2.59$	$\times 4.39$	$\times 4.27$	$\times \mathbf{6.26}$	$\times 5.00$

ing step is applied to complex-valued fMRI images. However, when considering magnitude-only images as input parameters to denoising, Hybrid-PCA on average is the best performer, whether we look at the whole brain or ROI-restricted statistical analysis. This demonstrates that the behavior of this approach reported for subject 3 in Section 4.1 is not fully representative of the 6 participants.

4.3. Noise map estimation

MP-PCA and Hybrid-PCA give access to an a posteriori estimate of the noise variance map (not shown). Firstly, this extra information can be used by NORDIC [15], which takes as optional input parameter a g -factor map. The latter is usually cumbersome to compute as it necessitates to collect additional noise-only data. However, once computed, this map can be used to perform a spatial normalization.

Secondly, when considering MP-PCA or OT approaches that do not rely on this extra knowledge by default, this noise map could be further used in the statistical model for fMRI data analysis to solve a weighted ordinary least squares (OLS) when fitting the GLM parameters instead of a single OLS.

5. CONCLUSION

In this study, we tested multiple LLR denoising methods as an additional step of high resolution fMRI images preprocessing. We demonstrated the added value of injecting this step as early as possible in the workflow in terms of gain in statistical sensitivity. Among the 5 tested methods, we found that the Optimal Threshold approach provides the best sensitivity-specificity trade-off as it does not retrieve false positives in the white matter while boosting the statistical significance in the visual cortex. Further we also realized there is not a *one size fits all solutions*: While the OT method yields the best results when performing the denoising step on complex-valued fMRI images (which are not always available), Hybrid-PCA is the most promising when denoising magnitude-only images. This preliminary study on task-based fMRI calls for a broader validation, both on a larger cohort and in other use cases, notably in resting-state fMRI. Importantly, the use of extra data such as a noise map in Hybrid-PCA and Hybrid-OT is not mandatory to already obtain a significant improvement: Sometimes it helps (Hybrid-PCA vs MP-PCA), sometimes it does not (Hybrid-OT vs OT). Future work also involves analyzing the impact of hyperparameters of denoising methods (e.g. the patch size and their overlap) on the downstream task.

6. COMPLIANCE WITH ETHICAL STANDARDS

The in vivo experimental protocol was approved by the local and national ethical committees the latest was filed under the identifier CPP10048 and issued by the National Comité de Protection des Personnes (CPP Sud Méditerranée 4 number 180913, IDRCB: 2018-A0011761-53). All participants gave their informed consent.

References

- [1] S. A. Huettel, S. Allen W, and M. Gregory, *Functional Magnetic Resonance Imaging*. Sunderland, Mass.: Sinauer Associates, Publishers, 2004.
- [2] F. G. Ashby, *Statistical Analysis of fMRI Data, Second Edition*. MIT Press, Sep. 17, 2019, 569 pp. Google Books: [yE2rDwAAQBAJ](http://books.google.com/books?id=yE2rDwAAQBAJ).
- [3] S. M. Smith, M. Jenkinson, M. W. Woolrich, *et al.*, “Advances in functional and structural MR image analysis and implementation as FSL,” *NeuroImage*, vol. 23 Suppl 1, S208–219, 2004. PMID: 15501092.
- [4] Z. Amor, G. Chaithya, G. Daval-Fr  rot, *et al.*, “Prospects of non-Cartesian 3D-SPARKLING encoding for functional MRI: A preliminary case: Study for retinotopic mapping,” *Proceedings of the 30th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, 2022.
- [5] M. Barth and B. A. Poser, “Advances in High-Field BOLD fMRI,” *Materials*, vol. 4, no. 11, pp. 1941–1955, Nov. 2, 2011.
- [6] C. Triantafyllou, J. R. Polimeni, and L. L. Wald, “Physiological noise and signal-to-noise ratio in fMRI with multichannel array coils,” *NeuroImage*, vol. 55, no. 2, pp. 597–606, Mar. 15, 2011.
- [7] T. T. Liu, “Noise contributions to the fMRI signal: An overview,” *NeuroImage*, vol. 143, pp. 141–151, Dec. 1, 2016.
- [8] K. J. Friston, Ed., *Statistical Parametric Mapping: The Analysis of Functional Brain Images*, 1st ed. Amsterdam ; Boston: Elsevier/Academic Press, 2007, 647 pp.
- [9] O. Esteban, C. J. Markiewicz, R. W. Blair, *et al.*, “fMRIPrep: A robust preprocessing pipeline for functional MRI,” *Nature Methods*, vol. 16, no. 1, pp. 111–116, 1 Jan. 2019.
- [10] K. Gorgolewski, C. Burns, C. Madison, *et al.*, “Nipype: A Flexible, Lightweight and Extensible Neuroimaging Data Processing Framework in Python,” *Frontiers in Neuroinformatics*, vol. 5, 2011.
- [11] J. V. Manj  n, P. Coup  , L. Concha, A. Buades, D. L. Collins, and M. Robles, “Diffusion Weighted Image Denoising Using Overcomplete Local PCA,” *PLOS ONE*, vol. 8, no. 9, e73021, Sep. 3, 2013.
- [12] J. Veraart, D. S. Novikov, D. Christiaens, B. Ades-Aron, J. Sijbers, and E. Fieremans, “Denoising of diffusion MRI using random matrix theory,” *NeuroImage*, vol. 142, pp. 394–406, Nov. 15, 2016. PMID: 27523449.
- [13] R. Neto Henriques, S. N  rholm Jespersen, and N. Shemesh, “Hybrid PCA denoising - improving PCA denoising in the presence of spatial correlations,” *Proceedings of the 30th Scientific Meeting of the International Society for Magnetic Resonance in Medicine*, 2022.
- [14] L. Vizioli, S. Moeller, L. Dowdle, *et al.*, “Lowering the thermal noise barrier in functional brain mapping with magnetic resonance imaging,” *Nature Communications*, vol. 12, no. 1, p. 5181, 1 Aug. 30, 2021.
- [15] S. Moeller, P. K. Pisharady, S. Ramanna, *et al.*, “Noise reduction with DIstribution Corrected (NORDIC) PCA in dMRI with complex-valued parameter-free locally low-rank processing,” *NeuroImage*, vol. 226, p. 117 539, Feb. 1, 2021.
- [16] M. Gavish and D. L. Donoho. “The Optimal Hard Threshold for Singular Values is $4/\sqrt{3}$.” arXiv: [arXiv:1305.5870](https://arxiv.org/abs/1305.5870). (Jun. 4, 2014), [Online]. Available: <http://arxiv.org/abs/1305.5870> (visited on 06/23/2022), preprint.
- [17] M. Gavish and D. L. Donoho, “Optimal Shrinkage of Singular Values,” *IEEE Transactions on Information Theory*, vol. 63, no. 4, pp. 2137–2152, Apr. 2017.
- [18] J. Josse and S. Sardy. “Adaptive Shrinkage of singular values.” arXiv: [arXiv:1310.6602](https://arxiv.org/abs/1310.6602). (Nov. 22, 2014), [Online]. Available: <http://arxiv.org/abs/1310.6602> (visited on 09/22/2022), preprint.
- [19] A. Shabalin and A. Nobel. “Reconstruction of a Low-rank Matrix in the Presence of Gaussian Noise.” arXiv: [arXiv:1007.4148](https://arxiv.org/abs/1007.4148). (Jul. 23, 2010), [Online]. Available: <http://arxiv.org/abs/1007.4148> (visited on 09/22/2022), preprint.
- [20] H. Gudbjartsson and S. Patz, “The rician distribution of noisy mri data,” *Magnetic Resonance in Medicine*, vol. 34, no. 6, pp. 910–914, 1995.
- [21] E. J. Cand  s, C. A. Sing-Long, and J. D. Trzasko, “Unbiased Risk Estimates for Singular Value Thresholding and Spectral Estimators,” *IEEE Transactions on Signal Processing*, vol. 61, no. 19, pp. 4643–4657, Oct. 2013.
- [22] F. A. Breuer, S. A. Kannengiesser, M. Blaimer, N. Seiberlich, P. M. Jakob, and M. A. Griswold, “General formulation for quantitative G-factor calculation in GRAPPA reconstructions,” *Magnetic Resonance in Medicine*, vol. 62, no. 3, pp. 739–746, Sep. 2009.
- [23] V. A. Mar  enko and L. A. Pastur, “DISTRIBUTION OF EIGENVALUES FOR SOME SETS OF RANDOM MATRICES,” *Mathematics of the USSR-Sbornik*, vol. 1, no. 4, p. 457, Apr. 30, 1967.
- [24] 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, p. 715, Sep. 11, 2014. PMID: 25309391.
- [25] C. Bhushan, M. Chong, S. Choi, *et al.*, “Temporal Non-Local Means Filtering Reveals Real-Time Whole-Brain Cortical Interactions in Resting fMRI,” *PLoS ONE*, vol. 11, no. 7, e0158504, Jul. 8, 2016. PMID: 27391481.
- [26] *Public Protocols*, The Human Brain Project - brain charting task.
- [27] A. Abraham, F. Pedregosa, M. Eickenberg, *et al.*, “Machine learning for neuroimaging with scikit-learn,” *Frontiers in Neuroinformatics*, vol. 8, 2014.