

High-Quality 0.5mm Isotropic fMRI: Random Matrix Theory Meets Physics-Driven Deep Learning

Ömer Burak Demirel^{1,2}, Steen Moeller², Luca Vizioli², Burhaneddin Yaman^{1,2}, Logan Dowdle²,
Essa Yacoub², Kâmil Uğurbil², and Mehmet Akçakaya^{1,2}

Abstract—Submillimeter fMRI plays a vital role in studying the brain function at the mesoscale level, allowing investigation of functional activity in small cortical structures. However, such resolutions require extreme trade-offs between SNR, spatio-temporal resolution and coverage leading to numerous challenges. Therefore, interpretable locally low-rank denoising methods based on random matrix theory have been proposed and built into fMRI pipelines, but they require well-characterized noise distributions on reconstructed images, which hinders the use of emerging physics-driven deep learning reconstructions. In this work, we re-envision the conventional fMRI computational imaging pipeline to an alternative where denoising is performed prior to reconstruction. This allows for a synergistic combination of random matrix theory based thermal noise suppression and physics-driven deep learning reconstruction, enabling high-quality 0.5mm isotropic functional MRI. Our results show that the proposed strategy improves on denoising or physics-driven deep learning reconstruction alone, with better delineation of brain structures, higher tSNR particularly in mid-brain areas and the largest expected extent of activation in GLM-derived t-maps.

I. INTRODUCTION

Functional magnetic resonance imaging (fMRI) has revolutionized our understanding of the human brain by allowing non-invasive imaging of neural activity based on blood oxygenation level dependent (BOLD) signal [1], and has been a staple of numerous neuroscientific studies, including the Human Connectome Project (HCP) [2]. With recent advances in fMRI hardware, acquisition and processing, submillimeter resolution fMRI has received interest as a means for studying brain function at the mesoscale level [3]. At this scale, the most fundamental units of neural computations, such as layers and columns [4] can be examined. However, to adequately determine layer-specific activation maps through the columnar and laminar structures, fMRI acquisitions require ~ 0.5 mm isotropic or better resolution [5], [6]. Such target resolutions inherently cause diminished signal-to-noise ratio (SNR) and necessitate high acceleration rates for sufficient coverage of the brain.

Even at coarser resolutions, the poor SNR of fMRI acquisitions have made reducing the effects of thermal noise in fMRI image series a longstanding research topic. Early works for fMRI denoising relied on principal component

analysis (PCA) and independent component analysis (ICA) [7], [8] with heuristic strategies for selecting the retained components. The current state-of-the-art methods rely on a combination of locally low-rank (LLR) modeling and random matrix theory [6], [9]–[12]. Application of ideas from random matrix theory requires the noise in the image series to be well-characterized by a known distribution, e.g. Rician or Gaussian [13], [14]. Thus, they are typically used with linear reconstruction methods [15]–[17], such as parallel imaging, and are not compatible with non-linear reconstruction methods that lead to poorly understood reconstruction noise distributions [18].

Among such non-linear reconstruction methods, physics-driven deep learning (DL) has recently emerged as a powerful alternative approach for MRI reconstruction with improved image quality at higher acceleration rates [19]–[21]. In physics-driven DL reconstruction, MR physics knowledge is incorporated via the forward encoding operator to impose data fidelity with acquired measurements, while the regularization is learned implicitly via a neural network. Physics-driven deep learning has already been applied to fMRI successfully, showing that it does not alter the subsequent analysis [22], [23], for HCP-style acquisitions. However, physics-driven deep learning still faces challenges for very high-resolution fMRI, since at lower SNRs or very high accelerations, it may lead to over-regularization [24].

Thus, in order to enable high-quality fMRI at ~ 0.5 mm resolution, it may be desirable to combine the state-of-the-art in fMRI denoising and reconstruction, namely random matrix theory based LLR denoising methods with physics-driven DL techniques. However, the conventional fMRI computational imaging pipeline is built on performing reconstruction followed by denoising. As aforementioned, this makes the state-of-the-art fMRI denoising and MRI reconstruction methods incompatible. In this work, we propose a new workflow that overturns this traditional approach to enable denoising using random matrix theory based LLR methods prior to (non-linear) reconstruction. This allows for a synergistic combination of thermal noise suppression followed by DL reconstruction, leveraging the best of both worlds while unlocking high-quality 0.5mm isotropic resolution fMRI. The proposed strategy improves upon using noise suppression after reconstruction or physics-driven DL alone for 0.5mm isotropic fMRI acquisitions, both visually, and quantitatively in terms of tSNR and GLM-derived t-maps.

¹Department of Electrical and Computer Engineering, ²Center for Magnetic Resonance Research, University of Minnesota, MN, USA. First three authors contributed equally, e-mails: {demir035, moell018, lvizioli, yaman013, dowdl016, yaco0006, ugurb001, akcakaya}@umn.edu

II. METHODS

A. Locally Low Rank Model and Random Matrix Theory Based Denoising for fMRI

Denoising in the conventional computational MRI pipeline is performed on a reconstructed fMRI series. First, a $k_1 \times k_2 \times k_3$ patch of a volume from a time-frame $\tau \in \{1, \dots, T\}$ in the time series is extracted. This is then vectorized as $\mathbf{y}_\tau \in \mathbb{C}^M$, where $M = k_1 \cdot k_2 \cdot k_3$. Next, corresponding vectorized patches from all time-frames are concatenated to form a Casorati matrix $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_\tau, \dots, \mathbf{y}_T] \in \mathbb{C}^{M \times T}$. This noisy data across the time series is modeled as:

$$\mathbf{Y} = \mathbf{X} + \mathbf{N}, \quad (1)$$

where $\mathbf{X} \in \mathbb{C}^{M \times N}$ is the underlying data matrix, and $\mathbf{N} \in \mathbb{C}^{M \times N}$ is additive Gaussian noise. The LLR model assumes that the data matrix $\mathbf{X}$ for a given patch is low-rank [25], [26], and corresponding denoising methods typically perform singular value thresholding on the noisy data matrix $\mathbf{Y}$ to eliminate unwanted noise components (shown in Figure 1a) [10], [11], [27]. In particular, let the singular value decomposition of $\mathbf{Y}$ be $\mathbf{USV}^H$, where $\mathbf{S}$ is a diagonal matrix whose entries are the spectrum of ordered singular values $s_j, j \in \{1, \dots, \min(M, T)\}$. For a thresholding value of λ , the soft- [28] or hard-thresholded [29] matrix, $\mathbf{S}_\lambda$ is used to form the denoised matrix as $\mathbf{US}_\lambda \mathbf{V}^H$. The denoised images are then obtained by extracting the patches for each time-frame from these denoised matrices, and performing patch averaging to account for overlaps between patches [9], [11].

The state-of-the-art LLR approaches [9]–[12] use random matrix theory formulations to determine the aforementioned threshold λ automatically. Typically, these methods perform a number of processing steps to ensure that $\mathbf{N}$ has independent identically distributed (i.i.d.) entries. Subsequently, the threshold is determined by either asymptotic properties, such as the Marchenko-Pastur distribution on the singular values of $\mathbf{N}$ [11], or using non-asymptotic variants [9]. In

particular, in NOise Reduction with DIstribution Corrected (NORDIC) denoising [9], $\mathbf{N}$ being a Gaussian matrix was ensured by correcting for spatial noise variations in the reconstructed images due to g-factors [15], and threshold was determined using a non-asymptotic characterization of the largest singular value of such Gaussian matrices in a parameter-free manner. This leads to the interpretation that NORDIC removes all principal components of $\mathbf{Y}$ that cannot be distinguished from the Gaussian noise level. NORDIC applied after image reconstruction was shown to improve fMRI analysis without introducing image blurring across a variety of spatial resolutions and field strengths [6].

B. Physics-driven Deep Learning Reconstruction

The forward model for MRI acquisition is given as:

$$\mathbf{y}_\Omega = \mathbf{E}_\Omega \mathbf{x} + \mathbf{n}, \quad (2)$$

where $\mathbf{y}_\Omega \in \mathbb{C}^P$ are the acquired k-space measurements, $\mathbf{x} \in \mathbb{C}^{m_1 \times m_2}$ is the image of interest, $\mathbf{E}_\Omega : \mathbb{C}^{m_1 \times m_2} \rightarrow \mathbb{C}^P$ is the multi-coil encoding operator, Ω is the undersampling pattern, and $\mathbf{n}$ is the measurement noise. The regularized inverse problem for recovering $\mathbf{x}$ from $\mathbf{y}_\Omega$ is given as [30]:

$$\arg \min_{\mathbf{x}} \|\mathbf{y}_\Omega - \mathbf{E}_\Omega \mathbf{x}\|_2^2 + \mathcal{R}(\mathbf{x}), \quad (3)$$

where the quadratic term enforces data fidelity (DF) with the acquired k-space measurements and $\mathcal{R}(\cdot)$ is a regularizer. This objective function can be solved via variable splitting with quadratic penalty [31] that splits the optimization problem in (3) into two sub-problems:

$$\mathbf{z}^{(l-1)} = \arg \min_{\mathbf{z}} \mu \|\mathbf{x}^{(l-1)} - \mathbf{z}\|_2^2 + \mathcal{R}(\mathbf{z}), \quad (4)$$

$$\mathbf{x}^{(l)} = \arg \min_{\mathbf{x}} \|\mathbf{y}_\Omega - \mathbf{E}_\Omega \mathbf{x}\|_2^2 + \mu \|\mathbf{x} - \mathbf{z}^{(l-1)}\|_2^2, \quad (5)$$

where $\mathbf{x}^{(l)}$ is the reconstructed image at iteration l , $\mathbf{z}^{(l)}$ is an auxiliary image, and μ is the quadratic penalty parameter.

Physics-driven DL reconstruction alternates between Eq. (4)-(5) for a fixed number of iterations [18] in a process




Fig. 1: a) Noise suppression after image reconstruction with LLR model and random matrix theory based threshold (conventional paradigm). Local patches are extracted from reconstructed images to form Casorati matrices. Singular value thresholding is performed using a random matrix theory based threshold that removes unwanted noise components. Lastly, patch averaging is performed to form the denoised image series. b) An example of rank sub-additivity of the Casorati matrices at an acceleration rate of $R = 3$. When uniform undersampling is performed, the small patches fold onto other patches preserving LLR properties. c) NORDIC denoising is performed on aliased images from individual channels of acquired k-space prior to reconstruction.

called algorithm unrolling [32]. Though the DF sub-problem in Eq. (5) has a closed-form solution, it can be efficiently solved by iterative linear methods, such as conjugate gradient [33]. On the other hand, the proximal operation in Eq. (4) is solved implicitly using neural networks [19], [20]. Finally, this unrolled network is trained end-to-end, typically using a loss function with respect to a reference image [18]. For such supervised training, this corresponds to the loss:

$$\min_{\theta} \frac{1}{N} \sum_{n=1}^N \mathcal{L}(\mathbf{x}_{\text{ref},n}, f(\mathbf{y}_{\Omega,n}, \mathbf{E}_{\Omega,n}; \theta)), \quad (6)$$

where $f(\mathbf{y}_{\Omega}, \mathbf{E}_{\Omega}; \theta)$ is the output of the unrolled network parametrized by θ , N is the number of training datasets in the database, $\mathcal{L}(\cdot, \cdot)$ is a loss function, and $\mathbf{x}_{\text{ref},n}$, $\mathbf{y}_{\Omega,n}$ and $\mathbf{E}_{\Omega,n}$ are the reference image, acquired k-space data and the corresponding multi-coil encoding operator for the n^{th} training sample.

III. PROPOSED COMPUTATIONAL IMAGING PIPELINE FOR HIGH-RESOLUTION FMRI

While denoising methods based on the LLR model and random matrix theory based thresholds, and reconstruction approaches based on physics-driven deep learning offer state-of-the-art for their respective tasks, they are not compatible with each other in the conventional pipeline of reconstruction followed by denoising. Thus, our strategy is to re-envision this processing pipeline, and to enable the application of LLR denoising methods based on random matrix theory prior to physics-driven deep learning reconstruction.

When the data are acquired with uniformly undersampled patterns, as in fMRI with echo planar imaging (EPI), the aliasing artifact is a foldover artifact (shown in Figure 1b), in which R pixels of the full field-of-view image are folded onto each other in this aliased field-of-view, where R is the acceleration rate. The effect is similar on patches in the image, where a patch in the undersampled image with foldover artifacts corresponds to the summation of R patches from the full field-of-view image. Repeating the process across the fMRI time series, and noting the sub-additivity of matrix rank

$$\text{rank}\left(\sum_{r=1}^R \mathbf{Y}_r\right) \leq \sum_{r=1}^R \text{rank}(\mathbf{Y}_r), \quad (7)$$

we note that the Casorati matrices for the patches in the undersampled images are likely to be low-rank if the Casorati matrices corresponding to the patches in the full field-of-view image are sufficiently low-rank (Fig. 1c).

Based on this observation, we modify the NORDIC approach [6], [9] for pre-reconstruction denoising using several steps. First, since the aliasing artifacts occur independently in each receiver coil, we process the data from each coil individually, where the acquisition noise is i.i.d. in nature. This also eliminates the need for the g-factor correction in the original NORDIC method. Thus, the acquired undersampled k-space for a given coil is first converted to image domain, albeit with the foldover aliasing artifacts. Then,




Fig. 2: A schematic of self-supervised physics-driven deep learning reconstruction. Algorithm unrolling alternates between data fidelity (DF) and regularizer units. DF units incorporate the MRI physics information through the encoding operator using coil sensitivity maps and sampling masks, whereas regularizer is learned by a neural network. For self-supervised training without fully-sampled data, multi-mask SSDU approach is used, where the acquired k-space indices, Ω , are split into multiple pairs of disjoint sets Θ_k and Λ_k . Θ_k is used in DF units and Λ_k is to define training loss.

image patches from different time-frames in the fMRI image series are extracted, vectorized and concatenated to form the noisy and aliased Casorati matrices. This is followed by singular value thresholding, using an LLR model based on the sub-additivity of the matrix rank. Here, the threshold is again chosen based on the previous random matrix theory characterization [9]. This is followed by patch averaging to generate the denoised foldedover images for each time-series for the given coil. Finally, these images are taken back to undersampled k-space for each coil (Fig. 1d).

The denoised k-space data is then used to train a physics-driven reconstruction neural network. Since fully-sampled reference data is not available, the supervised training strategy described in Section II-B is not applicable. Instead, we rely on unsupervised strategies that allow training without fully-sampled data [34]. In particular, we use a self-supervised learning technique, called Self-Supervised learning via Data Undersampling (SSDU) [21], [35], which splits the acquired k-space locations, Ω into two disjoint sets Θ and Λ . Θ is used in the DF units while Λ remains unseen by the network and is used to define the k-space training loss. To further improve the performance, multiple disjoint pairs of (Θ_k, Λ_k) $k \in \{1, \dots, K\}$ can be used in a multi-mask version of SSDU [36], [37]. This leads to following training loss

$$\min_{\theta} \frac{1}{N \cdot K} \sum_{n=1}^N \sum_{k=1}^K \mathcal{L}(\mathbf{y}_{\Lambda_k,n}, \mathbf{E}_{\Lambda_k,n}(f(\mathbf{y}_{\Theta_k,n}, \mathbf{E}_{\Theta_k,n}; \theta))), \quad (8)$$

where N is the number of training data in the database, $\mathcal{L}(\cdot, \cdot)$ is the loss function, and Θ_k, n and Λ_k, n are the k^{th} DF and loss masks for the n^{th} training data sample, respectively (Figure 2).

IV. EXPERIMENTS AND RESULTS

A. fMRI data acquisition

Imaging experiments were performed at 7T (Siemens Magnetom) in 3 subjects using a 32-channel NOVA head coil. The study was approved by our institutional review board and written informed consent was acquired before each scan. A T_2^* -weighted 3D GE-EPI sequence was performed that covers 40 slices with $TR = 83\text{ms}$ (Volume Acquisition Time = 3654 ms with 10% slice oversampling). The relevant imaging parameters were: $TE = 32.4\text{ms}$, flip angle = 13° , bandwidth = 820Hz, phase-encoding acceleration $R = 3$, partial-Fourier = 6/8, 0.5mm isotropic resolution. In total 8 runs were acquired each lasting around 5:30 minutes. All runs are collected with a standard 24s on, 24s off visual block design paradigm with a center/target and surrounded checkerboard counter phase flickering at 6 Hz.

B. Implementation details

NORDIC denoising was applied on the undersampled 3D-EPI images as described in Section III. After read-out oversampling was removed, eddy current and timing correction were applied. Then, an inverse Fourier transform was applied along each of the three k-space dimensions firstly each channel individually. A 3D spatial patch, with a spatial:temporal ratio of 11:1 was used. For an acquisition with $T \approx 90$, this corresponds to $10 \times 10 \times 10$ patches in the images with foldover along the phase-encoding direction. These were used to form the Casorati matrices for LLR modeling. For each channel, the thermal noise level was determined from the readout direction using the standard deviation of all the signals with the highest and lowest frequency. Using the standard deviation of the thermal noise, a matrix with i.i.d. entries of identical dimension to the Casorati matrix and identical standard deviation to the thermal noise was generated. From the i.i.d. generated matrix, the sample mean for the highest singular value was determined, and used as the threshold for singular value thresholding on the noisy Casorati matrices.

Following denoising, a physics-driven deep learning reconstruction network was trained with multi-mask SSDU [36] using $K = 3$ masks. First, the 3D-EPI k-space was inverse Fourier transformed along the slice direction, and these slices were processed individually leading less memory requirements. The physics-driven deep learning was unrolled for 10 iterations alternating between the regularizer and the DF sub-problems in Eq. (4)-(5). The latter was solved using conjugate gradient, which itself was unrolled for 10 iterations [33]. The proximal operator for the regularizer in (4) was solved by a convolutional neural network based on a ResNet structure [21]. Sensitivity maps are estimated using ESPIRiT [38] from a low resolution scan, and were used in DF units. A normalized ℓ_1 - ℓ_2 loss was used for $\mathcal{L}(\cdot, \cdot)$ [21]. Adam optimizer with learning rate of $3 \cdot 10^{-4}$ was used over 100 epochs. Training was performed using a total number of 352 2D k-spaces consisting of 2 subjects, each having 4 runs and 44 slices with 1 time-frame per subject.

Testing was performed on a different subject unseen by the network, where all runs, all slices and all time-frames were reconstructed. Note that during the deep learning reconstruction each time-frame was reconstructed individually, thus no temporal information was shared across image series.

Comparisons were made between four methods: 1) Conventional parallel imaging, using GRAPPA [16], performed on acquired (non-denoised) raw k-space, referred to as “Non-denoised GRAPPA”. 2) NORDIC denoising applied to k-space prior to GRAPPA reconstruction, referred to as “NORDIC-denoised GRAPPA”. 3) Physics-driven deep learning reconstruction performed on acquired (non-denoised) raw k-space, referred to as “Non-denoised Physics-Driven (PD) DL”. 4) The proposed method, where physics-driven deep learning reconstruction was performed on the NORDIC-denoised raw k-space, referred to as “NORDIC-denoised Physics-Driven (PD) DL”. GRAPPA kernels were calibrated using 5×4 kernel size for in-plane unaliasing, using the same calibration data utilized for the generation of the ESPIRiT coil maps for physics-driven deep learning. Separate trainings for non-denoised and NORDIC-denoised raw k-spaces were performed for deep learning reconstruction using the same setup, since the performance of deep learning methods have been reported to depend on the underlying SNR [39].

C. Data Analysis

Functional pre-processing was performed in BrainVoyager [40]. First, 3D rigid body motion correction was applied, where for each run, each volume was realigned to first volume of the first run using sinc interpolation. Additionally, de-trending was performed by regressing out low-drifts, up to 3rd order discrete cosine transform from the motion corrected time series. Subsequently, standard general linear model (GLM) with ordinary least squares minimization was performed to estimate BOLD-evoked response amplitudes. GLM design matrices were generated by convolving of a double gamma with a “box car” function, the latter representing the stimuli’s onsets and offsets. GLM analyses were performed on all runs concatenated for each reconstruction independently. For each voxel, percent signal change amplitudes were computed by dividing the GLM beta weights (representing BOLD evoked responses) by the mean of the pre-processed time series. Temporal SNR (tSNR) was computed on a pixel basis, by dividing the mean of the pre-processed time courses by their standard deviations.

D. Evaluation

Representative reconstructed slices are shown in Figure 3. Among all methods, non-denoised GRAPPA shows substantial amount of noise amplification rendering unusable image quality. NORDIC-denoised GRAPPA (1st row, 2nd column) and non-denoised physics-driven DL (2nd row, 1st column) both reduce the noise compared to non-denoised GRAPPA. Note some loss of details are seen in non-denoised physics-driven DL reconstruction shown by yellow arrows, indicative of spatial smoothing. The proposed NORDIC-denoised




Fig. 3: A representative slice of 0.5mm isotropic fMRI data with in-plane acceleration rate of 3. NORDIC-denoised GRAPPA reduces noise compared to non-denoised GRAPPA, which shows substantial noise amplification. Non-denoised physics-driven DL shows improved image quality compared to non-denoised GRAPPA, but shows loss of details compared to NORDIC-denoised GRAPPA (yellow arrows). NORDIC-denoised physics-driven DL reconstruction shows the highest image quality among all methods preserving the sharpness.

physics-driven DL reconstruction (2nd, 2nd column) shows visually the best image quality with reduced noise and preservation of fine details.

tSNR maps are depicted in Figure 4 for all methods. Non-denoised GRAPPA shows the lowest tSNR among all methods, while NORDIC-denoised GRAPPA substantially improves upon it. Although tSNR gain is also seen in the non-denoised physics-driven DL, there is lower tSNR in brain periphery regions compared to NORDIC-denoised GRAPPA. Additionally, these tSNR maps show anatomical structures, indicative of over-regularization in the non-denoised physics-driven DL. The proposed NORDIC-denoised physics-driven DL reconstruction shows the highest tSNR gain among all including gains in the central brain regions, with no discernible over-regularization.

Figure 5 shows GLM-derived t-maps for the contrast target and surround > 0 for all reconstructions. Non-denoised GRAPPA t-maps are dominated by thermal noise, leading to no meaningful activation. NORDIC-denoised GRAPPA and non-denoised physics-driven DL allow retrieval of the retinotopically expected extent of activation. The NORDIC-denoised physics-driven DL leads to the largest expected extent of activation.

V. DISCUSSION AND CONCLUSION

This paper proposes a new computational imaging pipeline for high-resolution fMRI to enable target voxel volumes of $\leq 0.1\mu\text{L}$ set forth by the Brain Initiative working group [41]. Particularly for layer fMRI, even 0.8mm resolution can be insufficient to capture subtle modulations across the 6 layers spanning a cortical thickness of 1.5mm to 4.5mm approximately, necessitating 0.5mm or higher resolutions [6]. By redesigning the standard pipeline of image reconstruction




Fig. 4: tSNR maps of a slice for all four methods. Non-denoised GRAPPA shows the lowest tSNR. NORDIC-denoised GRAPPA and non-denoised physics-driven DL reconstruction improve upon it substantially, with similar tSNR levels. However, non-denoised physics-driven DL shows anatomical structures, indicative of over-regularization. The proposed NORDIC-denoised physics-driven DL reconstruction shows the highest tSNR among all methods, including substantial gains in central brain regions.

followed by denoising, our work enables a synergistic combination of state-of-the-art fMRI denoising methods based on LLR modeling and random matrix theory with physics-driven deep learning reconstruction. Our results show that the proposed processing outperforms physics-driven deep learning or NORDIC denoising alone, both visually, and in terms of tSNR and GLM-derived t-maps, enabling high-quality 0.5mm isotropic resolution fMRI.

This study has limitations. Although multiple runs were




Fig. 5: GLM-derived t-maps for the contrast target using four different reconstructions. Non-denoised GRAPPA is dominated by thermal noise and does not show any meaningful activation. NORDIC-denoised GRAPPA and non-denoised physics-driven DL reveal retinotopically expected extent of activations. NORDIC-denoised physics-driven DL reconstruction (2nd, 2nd column) shows the largest expected extent of activation.

acquired for each subject, leading to higher statistical power [6], further evaluation in a larger cohort is warranted. This work exploited the retinotopic properties of V1, which provide a ground truth as to the spatial extent of activation elicited by our stimuli (or the flickering gratings) [6]. The performance of the neural network for the physics-driven deep learning reconstruction may further be improved with higher capacity neural networks, but this may require more training samples. Additionally, further investigations at higher acceleration rates are also warranted to harness the full potential at higher spatial resolutions with the proposed combination.

ACKNOWLEDGMENT

This work was partially supported by NIH: P41 EB027061, R01 HL153146, U01 EB025144, P30 NS076408, RF1 MH116978; NSF: CAREER CCF-1651825.

REFERENCES

- [1] N. K. Logothetis, et al., "Neurophysiological investigation of the basis of the fMRI signal," *Nature*, vol. 412, no. 6843, pp. 150–157, 2001.
- [2] M. F. Glasser, et al., "The human connectome project's neuroimaging approach," *Nature Neuroscience*, vol. 19, no. 9, pp. 1175–1187, 2016.
- [3] M. Fukuda, et al., "Submillimeter-resolution fMRI: Toward understanding local neural processing," *Progress in Brain Research*, vol. 225, pp. 123–152, 2016.
- [4] E. Yacoub, et al., "High-field fMRI unveils orientation columns in humans," *Proceedings of the National Academy of Sciences*, vol. 105, no. 30, pp. 10607–10612, 2008.
- [5] E. S. Finn, et al., "Higher and deeper: Bringing layer fMRI to association cortex," *Progress in Neurobiology*, vol. 207, pp. 101930, 2021.
- [6] L. Vizioli, et al., "Lowering the thermal noise barrier in functional brain mapping with magnetic resonance imaging," *Nature Communications*, vol. 12, no. 1, pp. 1–15, 2021.
- [7] P. P. Mitra, et al., "The nature of spatiotemporal changes in cerebral hemodynamics as manifested in functional magnetic resonance imaging," *Magnetic Resonance in Medicine*, vol. 37, no. 4, pp. 511–518, 1997.
- [8] C. G. Thomas, et al., "Noise reduction in BOLD-based fMRI using component analysis," *Neuroimage*, vol. 17, no. 3, pp. 1521–1537, 2002.
- [9] S. Moeller, et al., "NOise reduction with DIstribution Corrected (NORDIC) PCA in dMRI with complex-valued parameter-free locally low-rank processing," *NeuroImage*, vol. 226, pp. 117539, 2021.
- [10] J. Veraart, et al., "Diffusion MRI noise mapping using random matrix theory," *Magnetic Resonance in Medicine*, vol. 76, no. 5, pp. 1582–1593, 2016.
- [11] J. Veraart, et al., "Denoising of diffusion MRI using random matrix theory," *Neuroimage*, vol. 142, pp. 394–406, 2016.
- [12] L. Cordero-Grande, et al., "Complex diffusion-weighted image estimation via matrix recovery under general noise models," *Neuroimage*, vol. 200, pp. 391–404, 2019.
- [13] A. M. Wink and J. B. Roerdink, "Bold noise assumptions in fMRI," *International Journal of Biomedical Imaging*, vol. 2006, 2006.
- [14] S. N. Sotiropoulos, et al., "Effects of image reconstruction on fiber orientation mapping from multichannel diffusion MRI: reducing the noise floor using SENSE," *Magnetic Resonance in Medicine*, vol. 70, no. 6, pp. 1682–1689, 2013.
- [15] K. P. Pruessmann, et al., "SENSE: sensitivity encoding for fast MRI," *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999.
- [16] M. A. Griswold, et al., "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.
- [17] S. F. Cauley, et al., "Interslice leakage artifact reduction technique for simultaneous multislice acquisitions," *Magnetic resonance in medicine*, vol. 72, no. 1, pp. 93–102, 2014.
- [18] F. Knoll, et al., "Deep-learning methods for parallel magnetic resonance imaging reconstruction: A survey of the current approaches, trends, and issues," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 128–140, 2020.
- [19] J. Schlemper, et al., "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE Transactions on Medical Imaging*, vol. 37, no. 2, pp. 491–503, 2017.
- [20] K. Hammernik, et al., "Learning a variational network for reconstruction of accelerated MRI data," *Magnetic Resonance in Medicine*, vol. 79, no. 6, pp. 3055–3071, 2018.
- [21] B. Yaman, et al., "Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data," *Magnetic Resonance in Medicine*, vol. 84, no. 6, pp. 3172–3191, 2020.
- [22] O. B. Demirel, et al., "Improved simultaneous multi-slice functional MRI using self-supervised deep learning," in *2021 55th Asilomar Conference on Signals, Systems, and Computers*, 2021, pp. 890–894.
- [23] O. B. Demirel, et al., "20-fold accelerated 7T fMRI using referenceless self-supervised deep learning reconstruction," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2021, pp. 3765–3769.
- [24] M. J. Muckley, et al., "Results of the 2020 fastMRI challenge for machine learning MR image reconstruction," *IEEE transactions on medical imaging*, vol. 40, no. 9, pp. 2306–2317, 2021.
- [25] J. P. Haldar and Z.-P. Liang, "Low-rank approximations for dynamic imaging," in *2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*. IEEE, 2011, pp. 1052–1055.
- [26] J. Trzasko, et al., "Local versus global low-rank promotion in dynamic MRI series reconstruction," in *Proceedings of International Society for Magnetic Resonance in Medicine*, 2011, vol. 19, p. 4371.
- [27] C.-A. Deledalle, et al., "Image denoising with patch based PCA: local versus global," in *BMVC*, 2011, vol. 81, pp. 425–455.
- [28] D. L. Donoho, "De-noising by soft-thresholding," *IEEE Transactions on Information Theory*, vol. 41, no. 3, pp. 613–627, 1995.
- [29] D. Goldfarb and S. Ma, "Convergence of fixed-point continuation algorithms for matrix rank minimization," *Foundations of Computational Mathematics*, vol. 11, no. 2, pp. 183–210, 2011.
- [30] M. Lustig, et al., "Sparse mri: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [31] J. A. Fessler, "Optimization methods for magnetic resonance image reconstruction: Key models and optimization algorithms," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 33–40, 2020.
- [32] K. Gregor and Y. LeCun, "Learning fast approximations of sparse coding," in *Proceedings of the 27th International Conference on Machine Learning*, 2010, pp. 399–406.
- [33] H. K. Aggarwal, et al., "MoDL: Model-based deep learning architecture for inverse problems," *IEEE Transactions on Medical Imaging*, vol. 38, no. 2, pp. 394–405, 2018.
- [34] M. Akçakaya, et al., "Unsupervised deep learning methods for biological image reconstruction and enhancement: An overview from a signal processing perspective," *IEEE Signal Processing Magazine*, vol. 39, no. 2, pp. 28–44, 2022.
- [35] B. Yaman, et al., "Self-supervised physics-based deep learning MRI reconstruction without fully-sampled data," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2020, pp. 921–925.
- [36] B. Yaman, et al., "Multi-mask self-supervised learning for physics-guided neural networks in highly accelerated magnetic resonance imaging," *NMR Biomed*. 35(12):e4798, 2022.
- [37] B. Yaman, et al., "Ground-truth free multi-mask self-supervised physics-guided deep learning in highly accelerated MRI," in *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2021.
- [38] M. Uecker, et al., "ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA," *Magn Reson Med*, vol. 71, no. 3, pp. 990–1001, 2014.
- [39] F. Knoll, et al., "Assessment of the generalization of learned image reconstruction and the potential for transfer learning," *Magnetic Resonance in Medicine*, vol. 81, no. 1, pp. 116–128, 2019.
- [40] R. Goebel, "BrainVoyager—past, present, future," *Neuroimage*, vol. 62, no. 2, pp. 748–756, 2012.
- [41] C. Bargmann, et al., "Brain 2025: a scientific vision," *Brain Research through Advancing Innovative Neurotechnologies (BRAIN) Working Group Report to the Advisory Committee to the Director, NIH*, 2014.