

PLUG-AND-PLAY RECONSTRUCTION FOR 3D NON-CARTESIAN fMRI DATA

Pierre-Antoine Comby^{*†}

Matthieu Terris[†]

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) captures a temporal sequence of 3D T_2^* weighted magnetic resonance volumes to study brain activity at rest or during task performance. Consequently, imaging at the required high temporal and/or spatial resolution necessitates highly undersampled k-space acquisitions, making the reconstruction of these volumes severely ill-posed. Yet, due to the lack of ground truth in brain activity, evaluating the impact of reconstruction methods on the statistical analysis of fMRI remains an open problem. In this work, we leverage the SNAKE-fMRI simulator to generate realistic artificial brain activity over a 3D brain template and synthesize the corresponding k-space data, providing a reference ground truth for evaluating reconstruction methods in the fMRI statistical analysis stage. Additionally, we introduce a Plug-and-Play (PnP) approach for image reconstruction that relies on an implicit 3D deep denoising prior. In a benchmark against compressed sensing strategies, we show that PnP-based reconstruction strategies produce high-quality images but suffer from a loss in statistical sensitivity compared to compressed sensing reconstructions in simulated fMRI studies when a ground truth is available.

Index Terms— functional MRI; Plug-and-Play; Compressed Sensing

1. INTRODUCTION

Functional Magnetic Resonance Imaging (fMRI) is a key neuroimaging technique to probe brain activity non-invasively. It relies on the blood oxygen level-dependent (BOLD) contrast as an indirect measure of neural activity due to neurovascular coupling [1]. fMRI data consists of a sequence of 3D T_2^* weighted volumes captured with spatial resolution at the millimeter level and second-level temporal resolution. The BOLD signal is sensitive to changes in blood oxygenation, which standard MRI scanners can detect. More specifically, as neurons activate and consume more energy and oxygen, the vascular system provides local fresh blood in return with a higher concentration of oxyhemoglobin, which can be detected through an extended relaxation parameter T_2^* and thus a transient increase in the magnitude of the MRI signal.

fMRI data acquisition consists of several steps. First, raw data are acquired in the spatial Fourier domain (k-space) while the volunteer is lying on the scanner, either performing some task or staying still at rest. These data are then reconstructed into a temporal sequence of 3D brain volumes (often called images), followed by pre-processing steps to correct for artifacts such as motion and off-resonance effects (blurring artifacts and signal loss) due to B0 inhomogeneities. Finally, statistical analysis is performed to detect signal variations associated with the experimental paradigm, enabling the investigation of neural correlates of specific cognitive functions.


Ensuring high-quality processing at each stage is, in turn, crucial to obtaining reliable statistical inferences on neural activity.

During recent decades, the fMRI pipeline has continuously evolved, benefiting from advances in anatomical MRI acquisition and reconstruction, such as parallel imaging (Cartesian acquisitions) [2], compressed sensing (non-Cartesian acquisitions) [3]–[5], and more recently, deep learning for MRI reconstruction [6]–[9]. These techniques have significantly allowed MR physicists to reach a higher spatial resolution in decent scan times through more prominent undersampling factors in k-space, mainly using 3D non-Cartesian trajectories (e.g., 3D SPARKLING)[10]–[12] to ensure isotropic resolution. However, the larger the undersampling, the more ill-posed the image reconstruction problem.


In fMRI, as time becomes an additional dimension in data acquisition and image reconstruction, strategies must be carefully tailored to balance spatial resolution, temporal resolution, signal-to-noise ratio (SNR), and brain coverage. High SNR can be secured by performing scans at ultra-high magnetic field (7Tesla or above). However, the current race towards higher spatial and temporal resolution, ultimately to differentiate brain activity patterns across cortical layers (e.g., at 500 μ m resolution) [13], [14] or to uncover fast hemodynamic responses (e.g., at 0.5 Hz) [15], remains an unmet challenge for a complete brain coverage. An additional difficulty is the substantial variability of fMRI responses (evoked brain activity) across individuals, which impedes the optimization of this acquisition/reconstruction pipeline in vivo in healthy volunteers as a ground truth is missing.

We propose optimizing the data acquisition/image reconstruction pipeline *in silico* using the SNAKE simulator [16] to overcome this issue. SNAKE allows us to synthesize realistic artificial fMRI data under various acquisition setups and, hence, to select the best according to different metrics, namely image quality and the statistical sensitivity/specificity compromise, which are not necessarily aligned. In this work, we introduce a novel learning-based prior, based on a 3D implicit denoiser, injected into a Plug-and-Play (PnP) method for fMRI image reconstruction and compare the performance of deep learning and variational methods for image reconstruction on both image quality and statistical criteria.

The remainder of this paper is structured as follows. First, we describe the simulation setup to produce realistic and accelerated artificial task-related fMRI data in k-space. Then we provide some background on PnP methods based on denoising priors and justify our choice of neural network architecture. Next, we introduce our benchmark that compares zero-filled adjoint Fourier, Compressed Sensing, and deep learning for fMRI image reconstruction on a frame-by-frame basis. Finally, we present and discuss our benchmark results by assessing each approach’s ability to recover decent image quality and good statistical performances to detect evoked brain activity.



(a) Brain Phantom on a field of view of (192,192,144) mm with the T_2^* -weighted contrast prescribed by TR, TE, FA . Activated region in the occipital cortex in red.



(b) Example of a block design paradigm for fMRI.

Fig. 1: Brain Phantom used in our simulation (top) and average BOLD response (simulated and measured) in the ROI (bottom).

2. MATERIALS AND METHODS

2.1. Acquisition simulation setup

To generate realistic artificial fMRI data in both the time-image and k-space domains, we use the SNAKE simulator [16]. We first consider a brain phantom, depicted in Fig. 1(a), with a true T_2^* -weighted contrast. To achieve this, we simulate a 3D Gradient Recall Echo acquisition (repetition time $TR = 50$ ms, echo time $TE = 22$ ms, flip angle $FA = 12^\circ$) using an 8-channel phased array coil ($L = 8$). Next, we simulate BOLD activity according to the block design paradigm (black trace in Fig. 1(b)), alternating 20 seconds-long blocks of visual stimulation and rest in the primary visual cortex based on the Harvard-Oxford atlas. The brain region eliciting the evoked activity is delineated by the red contour in Fig. 1(a). The magnitude of the BOLD effect is defined by multiplying the fuzzy gray matter tissue mask voxel-wise with the mask of the primary visual cortex. In each voxel, we then generate the simulated BOLD time series by scaling the block-design regressor (black trace in Fig. 1(b)) with the magnitude of the BOLD effect and convolving the result with a canonical hemodynamic response, resulting in the orange trace in Fig. 1(b). The stimulation lasts $T = 300$ sec in total (one half shown in Fig. 1b).

To simulate a realistic MR acquisition process, we consider the presence of N_{tis} tissue classes in the 3D brain phantom. This means that we model tissue-specific parameters, namely tissue proportion $(w_i(\mathbf{r}))_{i=1}^{N_{\text{tis}}}$ per voxel \mathbf{r} , contrast $((\mu_i)_{i=1}^{N_{\text{tis}}})$ in the prescribed GRE sequence configuration, and transverse relaxation times $(T_{2,i}^*)_{i=1}^{N_{\text{tis}}}$ to accurately simulate a spatially varying MR imaging contrast, embodying T_2^* -decay. Then, regarding the simulation of fMRI data in k-space, we use a stack of in-out spirals as sampling trajectories $(\mathbf{k}_s(t))_{1 \leq s \leq n_{\text{shots}}}$ for n_{shots} consecutive shots of duration TR , at each timestep $0 \leq t \leq T$. The center of each spiral is placed at TE to optimize the T_2^* contrast. Since we start with a standard 3-mm isotropic resolution, full k-space coverage at the Nyquist rate typically con-




Fig. 2: Stack of spirals used to sample a single frame in 3D k-space. This pattern is repeated over time to collect the full fMRI sequence.

sists of 48 spirals, resulting in a volumetric TR (TR_v) of 2.4 sec. Here, as we are interested in high temporal resolution ($TR_v = 0.75$ seconds), we subsample the k-space data by a factor of 3.2 along the stacking dimension k_z (see Fig. 2), which consists of retaining only $n_f = 15$ k-space planes (3 central planes constantly sampled and 12 randomly picked up). As a consequence, this procedure produces an acceleration factor of 3.2 and $T/TR_v = 400$ fMRI scans. In this context, the simulated data $y_{\ell,s}(t)$ for the s -th shot on the ℓ -th coil follows an extended Fourier model:

$$y_{\ell,s}(t) = \int_{\text{FOV}} \sum_{i=1}^{N_{\text{tis}}} w_i(\mathbf{r}) \mu_i e^{-t/T_{2,i}^*} \mathcal{S}_\ell(\mathbf{r}) e^{-2i\pi \mathbf{k}_s(t) \cdot \mathbf{r}} d\mathbf{r} + \varepsilon_{\ell,s}(t), \quad (1)$$

where $\varepsilon_{\ell,s} \sim \mathcal{N}(0, \Sigma)$, $\forall \ell = 1, \dots, L$ is the thermal noise with homoscedastic coil covariance matrix Σ defined as follows:

$$\Sigma = \frac{1}{\text{SNR}} \left(\int_{\text{FOV}} \sum_{i=1}^{N_{\text{tis}}} w_i(\mathbf{r}) \mu_i d\mathbf{r} \right) \mathbf{I}_L, \quad (2)$$

where SNR refers to the signal-to-noise ratio. Importantly, the BOLD activation is updated at every unitary TR , modulating the contrast μ_i within a designated Region of Interest (ROI), referred to as the *Activation ROI tissue*. This region's intensity varies over time according to the prescribed BOLD effect for a given shot t , with a modulation amplitude of 2.5% relative to the baseline contrast ($\mu_{\text{ROI}} = 0.025 \cdot h(t)$, where $h(t)$ is the prescribed hemodynamic response). The maximum variation of $y_{\ell,s}(t)$ due to the BOLD effect is therefore impacted similarly. More details on the physics and computational aspects of the simulation process are available in [16].

Once all shots $\mathbf{y}_t = (y_{\ell,s}(t))_{s,\ell}$ are sequentially generated according to Eq. (1), each group of n_f consecutive shots is combined to reconstruct a single volume, forming a frame in the fMRI sequence. It should be noted that information about the different tissue types or relaxation parameters is not available during the reconstruction stage. Eventually, the acquisition model for a single sought image \mathbf{x}_t writes:

$$\mathbf{y}_t = \sum_{\ell=1}^L \mathcal{F}_{\Omega_t} \mathcal{S}_\ell \mathbf{x}_t + \varepsilon_{t,\ell} \quad (3)$$

where $\Omega_t = \{\mathbf{k}_s \mid tn_f \leq s \leq (t+1)n_f\}$ gathers all the shots of the t -th frame.

2.2. Reconstruction with PnP prior

Traditional methods for solving (3) often propose to estimate each \mathbf{x}_t by solving a regularized minimization problem of the form

$$\hat{\mathbf{x}}_t = \operatorname{argmin}_{\mathbf{x} \in \mathbb{C}^N} f(\mathbf{x}, \mathbf{y}_t) + \lambda r(\mathbf{x}), \quad (4)$$

where $f(\cdot, \mathbf{y}_t)$ is a data-fidelity enforcing functional (ensuring that the solution $\hat{\mathbf{x}}_t$ satisfies (3)), $r: \mathbb{C}^N \rightarrow \mathbb{R}$ ensures regularity of the solution, and $\lambda > 0$ is a regularization parameter. Here, f is chosen as $f(\mathbf{x}, \mathbf{y}_t) = \frac{1}{2} \sum_{\ell=1}^L \|F_{\Omega_t} \mathcal{S}_\ell \mathbf{x} - \mathbf{y}_{t,\ell}\|_2^2$, while r is chosen to promote the sparsity of the reconstruction in an appropriate domain. The minimizer $\hat{\mathbf{x}}_t$ is then computed with an optimization algorithm, typically involving proximity operators and/or gradients of f and r [17].

In this context, plug-and-play (PnP) algorithms propose to replace r in (3) with an implicitly defined regularization, typically through a deep neural network (DNN) trained as a denoiser. In practice, this operator replaces the proximity operator of r that appears in the minimization algorithm to solve (4). Following [18], we propose to perform the image reconstruction with a preconditioned Half-Quadratic Splitting (HQS) algorithm. More precisely, given a preconditioning matrix P and a number of iterations K , our reconstruction algorithm for estimating $\hat{\mathbf{x}}_t$ reads, for $0 \leq k \leq K$,

$$\begin{aligned} \mathbf{u}_{t,k} &= \operatorname{prox}_{\gamma f}^P(\mathbf{x}_{t,k}) \\ \mathbf{x}_{t,k+1} &= D_{\sigma_k}(\mathbf{u}_{t,k}), \end{aligned} \quad (\text{PnP-HQS})$$




where $\operatorname{prox}_{\gamma f}^P$ is the proximity operator of γf in the metric induced by P , $(\sigma_k)_{1 \leq k \leq K}$ is a decreasing sequence of noise levels, and where we set $\mathbf{x}_{t,0} = A^\dagger \mathbf{y}_t$. Interestingly, when $P = \text{Id}$, (PnP-HQS) reduces to the DPIR algorithm from [19] that has shown state-of-the-art performance in low-level vision tasks. In our experiments, we set $\gamma = 13$, $\sigma_k = \sigma_0 \xi^k$, with $\sigma_0 = 4 \cdot 10^{-3}$ and $\xi = 0.97$, and $K = 20$.

Here, we consider the ‘‘F-1’’ preconditioner $P = 2 - \alpha A^H A$ introduced in [20], where $A = F_{\Omega_t} \otimes \mathcal{S}_\ell$, and α is the Lipschitz constant of ∇f . Although other PnP algorithms may be chosen, two main observations motivate our choice: First, due to its annealing strategy of the parameters, it is fast and does not show the instabilities that other PnP algorithms may show [18]. Second, introducing the preconditioning matrix P ensures a robust reconstruction despite the high ill-conditioning of F_{Ω_t} in non uniformly sampled settings.

2.3. Neural network architecture

We next detail the chosen architecture for our denoiser D_σ . In this work, we consider complex volumetric data for $(\mathbf{x}_t)_{1 \leq t \leq T}$. More precisely, for all t , \mathbf{x}_t can be represented as a tensor of shape $(2, D, H, W)$, where the first two dimensions account for real and imaginary parts of the data. Consequently, the denoiser to be considered should be able to handle complex volumetric images. We therefore propose to extend the 2D bias-free DRUNet architecture from [19] into a 3D architecture by replacing the 2D convolution layers with the 3D convolutional layers, with 2 input channels. The resulting architecture enjoys the same theoretical properties as the original 2D architecture, *e.g.* its homogeneous property $D_{\alpha\sigma}(\alpha \mathbf{y}) = \alpha D_\sigma(\mathbf{y})$ for any $\alpha > 0$, ensuring good generalization abilities for this denoiser over images with various intensity ranges.

Table 1: Quantitative metrics on image quality and statistical performances. SSIM and PSNR scores are averaged over all frames. The balanced accuracy (BACC) is computed with a threshold set to $p < 10^{-3}$ and AUC and BACC were computed with respect to voxels with at least 50% of gray matter.

	Name	PSNR	SSIM	tSNR	AUC	BACC
	CG	18.560	0.367	nan	0.782	0.754
	CS	18.748	0.459	6.706	0.758	0.948
	HQS-F1	19.338	0.577	6.286	0.574	0.891

3. EXPERIMENTAL RESULTS

3.1. Benchmark of MR image reconstruction methods

We compare our reconstruction method with a wavelet-based compressed sensing reconstruction as implemented in `pysap-mri` [21]. More precisely, we use (4) with $r(\mathbf{x}) = \|\Psi \mathbf{x}\|_1$, where Ψ is a wavelet decomposition. Eq. (4) can be efficiently solved with a standard forward-backward algorithm [17]. We also provide comparisons with the pseudo-inverse $A^\dagger \mathbf{y}_t$, obtained by running a conjugate gradient (CG) algorithm on $f(\mathbf{x}, \mathbf{y}_t)$ for each frame t .

3.2. Training

We train our model on the Calgary-Campinas dataset [22]. We extract complex images by performing an inverse FFT on the raw Cartesian data in the k-space that will serve as our reference and combine multichannel inputs into a single 3D complex-valued volume using a virtual coil combination method [23]. We train our model using the Adam optimizer with a 10^{-4} learning rate. Pre-training is performed for 200k iterations on randomly extracted patches of size $64 \times 64 \times 64$ with a batch size of 12. The model is then fine-tuned in larger patches $128 \times 128 \times 128$ with a batch size of 2 for 100k iterations.

3.3. MR image reconstruction results

We now present the main results of our study. First, we compare the image quality of reconstructed fMRI sequences using CG, CS, and PnP with the ground truth using PSNR and SSIM scores in Tab. 1. In Fig. 3, we present the three reconstruction methods considered. The background image of each case corresponds to the first volume reconstructed in the fMRI sequence, and the selected cuts maximize the number of voxels displayed for the target ROI. The learning-based PnP-HQS reconstruction provides the best image quality. In contrast, the wavelet-based CS reconstruction provides lower image quality (0.12 lower SSIM scores) and embodies a few spiral aliasing artifacts. However, both regularized methods outperformed the classical CG reconstruction, which shows detrimental aliasing artifacts in the frontal and occipital regions.

3.4. Statistical analysis of fMRI sequences

Once reconstructed, we performed the regression analysis of fMRI signals using a simple general linear model (GLM) consisting of the regressor shown in orange in Fig. 1b and a constant baseline. The goal is to detect voxels eliciting evoked activity that correlates with the experimental paradigm. Then, z scores were calculated as the ratio of estimated regression weighted divided by the square root of the residuals in the GLM. Finally, z scores were thresholded at




Fig. 3: Z-score maps for CG/CS/HQS-F1 reconstruction methods (from left to right). the target ROI (voxel with at least 50% of gray matter, and belonging to primary visual cortex) is delineated in cyan. The complete anatomical ROI (occipital cortex) is contoured in green.




Fig. 4: Precision-Recall characteristic curve for the different thresholds of gray matter considered in the ROI.

a p-value of $p = 0.001$ uncorrected for multiple comparisons to assess each reconstruction pipeline’s statistical performances (sensitivity/specificity trade-off) and analyze to what extent the original ROI was recovered by computing scores and confusion metrics (true positive, false positive, etc.). As the ROI is small compared to the brain volume, we used precision and recall (true positive rate) to analyze the distribution of detected voxels.

Furthermore, voxels located in ROI have a spatially varying proportion of gray matter (GM), which directly impacts the magnitude of the BOLD effect in a proportional way. Thus, we considered different thresholds of z scores to define true positives. In Fig. 4, we present three different classification results for different GM proportions (voxels with respectively at least 10%, 50%, or 90%) in the respective dotted, dashed, and solid lines. Last, in Tab. 1, we compare the statistical performances for the CG/CS/HQS-F1 methods. As the CG method gives the highest AUC score while reporting the lowest image quality, this indicates a poor point spread function (PSF), which eventually blurs the true activations and produces false positives in addition to true positives. This is confirmed in Fig. 4 by lower precision and higher recall scores for CG reconstruction compared to CS and HQS-F1. In contrast, CS reconstruction reports the best AUC/BACC and highest Precision/Recall scores despite aliasing artifacts and poorer image quality compared to HQS-F1. Hence, CS achieves the best compromise between image quality and statistical performance. The lower performance of the learning-based PnP method might be due to the fact that the 3D DRUNet has been

trained on 1-mm isotropic images, while fMRI images were generated at 3-mm. Hence, additional fMRI simulations operating at 1-mm isotropic are necessary to confirm this finding soon, or instead to demonstrate the superiority of learning-based reconstruction methods, which have been trained at the target spatial resolution. However, because of their computational cost, these numerical experiments and the corresponding results were not available prior to submission.

4. CONCLUSION AND PERSPECTIVES

In this work, we studied how a learning-based prior embedded in a PnP method impacts image reconstruction in 3D non-Cartesian high-temporal fMRI. To do so, we used the SNAKE-fMRI simulator to draw preliminary conclusions in a fully controlled environment. Although image quality is significantly improved using the learning-based PnP method compared to standard techniques (CG, CS), our findings suggest a compromise between image quality and statistical performance in terms of true positives (recall)/false positive (precision) balance that is currently optimized by CS reconstruction. This could result from training a denoiser on a dataset of T_1 weighted MR images, while the target application is fMRI and involves a T_2^* weighed contrast. Furthermore, the resolution change between the training set and the test set could also explain the loss in statistical performance. we observed with PnP.

Future work will therefore focus on simulating scenarios at matched spatial resolution (i.e., 1-mm isotropic) and possibly on matching the contrasts between the training and test stages. This extension is computationally and memory challenging.

For the sake of reproducible science, detailed settings, implementations, and reproducible benchmark are available at <https://github.com/paquetteau/eusipco-pnp-fmri>. The study was carried out using the deepinv¹ and SNAKE-fMRI² libraries.

5. COMPLIANCE WITH ETHICAL STANDARDS

This is a numerical simulation study for which ethical approval was not required.

6. ACKNOWLEDGMENTS

This work was granted access to IDRIS’ HPC resources under the allocation 2023 AD011011153R3 made by GENCI. The authors have no relevant financial or non-financial interests to disclose.

¹<https://github.com/deepinv/deepinv>

²<https://github.com/mind-inria/snake-fmri>

References

- [1] S Ogawa, T. M. Lee, A. R. Kay, and D. W. Tank, "Brain magnetic resonance imaging with contrast dependent on blood oxygenation.," *Proceedings of the National Academy of Sciences*, vol. 87, no. 24, pp. 9868–9872, Dec. 1990.
- [2] M. A. Griswold, P. M. Jakob, R. M. Heidemann, *et al.*, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.
- [3] M. Lustig, D. Donoho, J. Santos, and J. Pauly, "Compressed Sensing MRI," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 72–82, Mar. 2008.
- [4] N. Chauffert, P. Ciuciu, J. Kahn, and P. Weiss, "Variable Density Sampling with Continuous Trajectories," *SIAM Journal on Imaging Sciences*, vol. 7, no. 4, pp. 1962–1992, Jan. 2014.
- [5] L. El Gueddari, C. Giliyar Radhakrishna, E. Chouzenoux, and P. Ciuciu, "Calibration-less multi-coil compressed sensing magnetic resonance image reconstruction based on oscar regularization," *Journal of Imaging*, vol. 7, no. 3, p. 58, 2021.
- [6] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model Based Deep Learning Architecture for Inverse Problems," *IEEE Transactions on Medical Imaging*, vol. 38, no. 2, pp. 394–405, Feb. 2019. eprint: 1712.02862 (cs).
- [7] K. Hammernik, T. Klatzer, E. Kobler, *et al.*, "Learning a variational network for reconstruction of accelerated MRI data," *Magnetic Resonance in Medicine*, vol. 79, no. 6, pp. 3055–3071, 2018.
- [8] M. J. Muckley, B. Riemenschneider, A. Radmanesh, *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, Sep. 2021.
- [9] Z. Ramzi, G. Chaithya, J.-L. Starck, and P. Ciuciu, "Nc-pdnet: A density-compensated unrolled network for 2d and 3d non-cartesian mri reconstruction," *IEEE Transactions on Medical Imaging*, vol. 41, no. 7, pp. 1625–1638, 2022.
- [10] C. Lazarus, P. Weiss, L. El Gueddari, *et al.*, "3d variable-density sparkling trajectories for high-resolution t2*-weighted magnetic resonance imaging," *NMR in Biomedicine*, vol. 33, no. 9, e4349, 2020.
- [11] G. R. Chaithya, P. Weiss, G. Daval-Fr  rot, A. Massire, A. Vignaud, and P. Ciuciu, "Optimizing Full 3D SPARKLING Trajectories for High-Resolution Magnetic Resonance Imaging," *IEEE Transactions on Medical Imaging*, vol. 41, no. 8, pp. 2105–2117, Aug. 2022.
- [12] Z. Amor, P. Ciuciu, C. GR, *et al.*, "Non-cartesian 3d-sparkling vs cartesian 3d-epi encoding schemes for functional magnetic resonance imaging at 7 tesla," *Plos one*, vol. 19, no. 5, e0299925, 2024.
- [13] J. R. Polimeni, V. Renvall, N. Zaretskaya, and B. Fischl, "Analysis strategies for high-resolution uhf-fMRI data," *Neuroimage*, vol. 168, pp. 296–320, 2018.
- [14] P. A. Bandettini, L. Huber, and E. S. Finn, "Challenges and opportunities of mesoscopic brain mapping with fMRI," *Current Opinion in Behavioral Sciences*, vol. 40, pp. 189–200, 2021.
- [15] O. Viessmann and J. R. Polimeni, "High-resolution fMRI at 7 Tesla: Challenges, promises and recent developments for individual-focused fMRI studies," *Current opinion in behavioral sciences*, vol. 40, pp. 96–104, Aug. 2021.
- [16] P.-A. Comby, A. Vignaud, and P. Ciuciu, "SNAKE-fMRI: A modular fMRI data simulator from the space-time domain to k-space and back." arXiv: 2404.08282 [eess]. (Apr. 12, 2024), [Online]. Available: <http://arxiv.org/abs/2404.08282>.
- [17] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," *Fixed-point algorithms for inverse problems in science and engineering*, pp. 185–212, 2011.
- [18] P.-A. Comby, B. Lapostolle, M. Terris, and P. Ciuciu, "Robust plug-and-play methods for highly accelerated non-cartesian MRI reconstruction," *ISBI 2025*, 2025.
- [19] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 6360–6376, 2021.
- [20] T. Hong, X. Xu, J. Hu, and J. A. Fessler, "Provable preconditioned plug-and-play approach for compressed sensing MRI reconstruction," *IEEE Transactions on Computational Imaging*, 2024.
- [21] S. Farrens, A. Grigis, L. El Gueddari, *et al.*, "Pysap: Python sparse data analysis package for multidisciplinary image processing," *Astronomy and Computing*, vol. 32, p. 100402, 2020.
- [22] R. Souza, O. Lucena, J. Garrafa, *et al.*, "An open, multi-vendor, multi-field-strength brain MR dataset and analysis of publicly available skull stripping methods agreement," *NeuroImage, Segmenting the Brain*, vol. 170, pp. 482–494, Apr. 15, 2018.
- [23] D. L. Parker, A. Payne, N. Todd, and J. R. Hadley, "Phase reconstruction from multiple coil data using a virtual reference coil," *Magnetic Resonance in Medicine*, vol. 72, no. 2, pp. 563–569, 2014.