



HAL
open science

Hybrid learning of Non-Cartesian k-space trajectory and MR image reconstruction networks

Chaithya Giliyar Radhakrishna, Zaccharie Ramzi, Philippe Ciuciu

► To cite this version:

Chaithya Giliyar Radhakrishna, Zaccharie Ramzi, Philippe Ciuciu. Hybrid learning of Non-Cartesian k-space trajectory and MR image reconstruction networks. ISBI 2022 - IEEE 19th International Symposium on Biomedical Imaging, IEEE, Mar 2022, Kolkata, India. hal-03394881

HAL Id: hal-03394881

<https://inria.hal.science/hal-03394881>

Submitted on 22 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HYBRID LEARNING OF NON-CARTESIAN K-SPACE TRAJECTORY AND MR IMAGE RECONSTRUCTION NETWORKS

Chaithya G R^{*†} Zaccharie RAMZI^{*‡} Philippe CIUCIU^{*†}

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

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

[‡]AIM, CEA, CNRS, Université Paris-Saclay, Université Paris Diderot

ABSTRACT

Compressed sensing (CS) in Magnetic resonance Imaging (MRI) essentially involves the optimization of 1) the sampling pattern in k-space under MR hardware constraints and 2) image reconstruction from the undersampled k-space data. Recently, deep learning methods have allowed the community to address both problems simultaneously, especially in the non-Cartesian acquisition setting. This paper aims to contribute to this field by tackling some major concerns in existing approaches. Regarding the learning of the sampling pattern, we perform ablation studies using parameter-free reconstructions like the density compensated (DCp) adjoint operator of the nonuniform fast Fourier transform (NUFFT) to ensure that the learned k-space trajectories actually sample the center of k-space densely. Additionally we optimize these trajectories by embedding a projected gradient descent algorithm over the hardware MR constraints. Later, we introduce a novel hybrid learning approach that operates across multiple resolutions to jointly optimize the reconstruction network and the k-space trajectory and present improved image reconstruction quality at 20-fold acceleration factor on T_1 and T_2 -weighted images on the fastMRI dataset with SSIM scores of nearly 0.92-0.95 in our retrospective studies.

1. INTRODUCTION

Speeding up magnetic resonance imaging (MRI) acquisitions involves jointly optimizing how the k-space is undersampled and how image quality is preserved (or aliasing artifacts are discarded) during image reconstruction. Based on the CS theory, the k-space has to be undersampled according to a variable density sampling (VDS) [1–5], whose shape depends on the underlying anatomy. In 2D imaging, 2D VDS patterns can only be efficiently achieved by using non-Cartesian sampling. Recently, SPARKLING [6–8] was introduced as an iterative procedure to optimize a non-Cartesian k-space sampling pattern according to a prescribed target sampling density (TSD) with each k-space trajectory being compliant with MR hardware constraints (particularly maximum gradient G_{\max} and slew rate S_{\max}). Further, the algorithm results in locally uniform sampling patterns and thus avoids holes and clusters in k-space. However, a major limitation of SPARKLING is the need for a TSD as an input in the optimization process. To address this issue, the TSD was learned in a deep learning setting in [9] using LOUPE [10] as an acquisition model. While this method provided adequate reconstruction performances, there was still a mismatch in the learning process. Indeed, the gridded TSD was learned in the Cartesian domain, while the actual trajectory being optimized was non-Cartesian. As this could lead to suboptimal results, there is a

need to jointly learn both the TSD and the reconstruction network in a non-Cartesian setting.

More recent methods [11–13] bypass the need for estimating a TSD, by jointly learning the trajectories and the reconstruction network in a data-driven manner. In [11, 13], the authors use a gridding step followed by a U-net to reconstruct the MR images and jointly learn the network parameters and k-space sampling trajectory. This method, however, relies on auto-differentiation of the NUFFT operator, which may not be very accurate (see [14]), thus resulting in suboptimal local minima. This was actually observed as the final learned trajectories only slightly deviated from their initialization. In [12], the authors use [14] to obtain a more accurate Jacobian approximation of the NUFFT operator. Both above referenced approaches [11, 12] enforce the hardware constraints with the help of a penalty on the cost function. Although a viable option, this does not guarantee that the optimized trajectories will meet these constraints. Additionally this requires tuning the hyperparameter associated with this additional penalty in the cost function. Further, this penalty could affect the overall gradients, thereby resulting in sub-optimality. To overcome these issues, the authors in [12] parameterized the trajectory with B-spline curves, which reduce the number of inequality constraints and hence the influence of this penalty. However, such parameterization could severely limit the trajectories and prevent them from better exploring the k-space. Finally, both methods do not make use of density compensation (DCp) which plays a crucial role in obtaining cleaner images in the non-Cartesian deep learning setting [15].

Given these results, in this work, we first present a generic model for jointly optimizing the non-Cartesian trajectories and the corresponding reconstruction network. More precisely, we introduce a method to learn the trajectories in a data-driven manner while embedding a projected gradient descent algorithm [16] to fulfill the hardware constraints. Second, we show experimentally that a novel generalized *hybrid learning* scheme performs better compared to a joint or alternated learning of the trajectory and the reconstruction network. Further we compare the trajectories obtained under different reconstruction settings (DCp Adjoint, DCp Adjoint + U-Net and NC-PDNet) and evaluate their performances in terms of image quality metrics (SSIM [17], PSNR) in a retrospective evaluation framework.

2. GENERIC MODEL

In this section we provide details (see Fig. 1) about the different components of the general model that is used for optimizing the non-Cartesian sampling trajectories as well as the image reconstruction network.

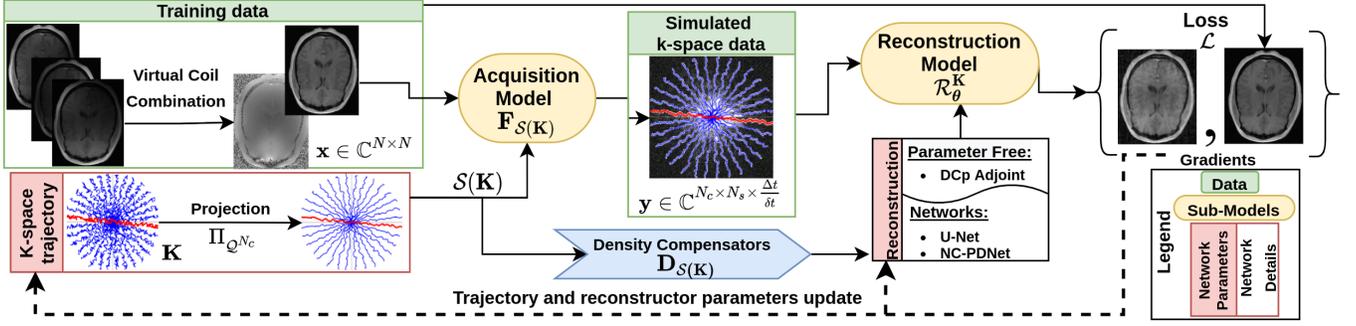


Fig. 1: A generic learning-based framework for joint optimization of the MRI acquisition and reconstruction models.

2.1. Acquisition

Data: In this work, we relied on the brain k-space data from the fastMRI dataset [18] for learning the k-space trajectories. We applied the virtual coil method [19] to combine per-channel images to obtain a complex-valued single channel image, which was used in our training stage. This was done to account for the phase accrual during k-space data acquisition and make the learned trajectories more realistic. Following the formulation developed in [8], we denote an MR image $\mathbf{x} \in \mathbb{C}^{N \times N}$, over a field of view $\mathcal{F} \times \mathcal{F}$. Then the 2D k-space of the image is defined in $[-K_{\max}, K_{\max}]^2$, with $K_{\max} = \frac{N}{2\mathcal{F}}$. In all our trajectories, we kept $N = 320$ and $\mathcal{F} = 0.23$ m. For the sake of simplicity, we normalize the k-space to $\Omega = [-1, 1]^2$.

Trajectory Specification: In our work, we directly learn the k-space trajectory locations unlike BJORK [12] where the coefficients of B-spline curves were learned. We believe that such parameterizations could potentially limit the exploration in the k-space. Instead we optimize the 2D k-space sampling pattern \mathbf{K} which is composed of N_c shots, $\mathbf{K} = (\mathbf{k}_i)_{i=1}^{N_c}$. Each shot can be played by the scanner hardware at the pace of gradient raster time Δt , throughout the readout time T_{obs} , resulting in $N_s = \lfloor \frac{T_{\text{obs}}}{\Delta t} \rfloor$ samples per shot. Thus the overall sampling pattern is $\mathbf{K} \in \Omega^{N_c \times N_s}$. Throughout this work, we used $N_s = 512$ ($T_{\text{obs}} \approx 5$ ms, for T_1 -w and T_2 -w imaging) and $N_c = 16$, leading to an acceleration factor $AF = \frac{N}{N_c} = 20$.

Forward Model: For the sake of simplicity, here we model the MR acquisition process using the NUFFT modeled by $\mathbf{F}_{\mathbf{K}}$ for the corresponding k-space sampling pattern \mathbf{K} . In practice, the k-space data is sampled at a δt -pace, with δt the dwell time used by the analog to digital converter (ADC). Hence a more realistic modeling of k-space data $\mathbf{y} \in \mathbb{C}^{N_c \times N_s \times \frac{\Delta t}{\delta t}}$ is as follows:

$$\mathbf{y} = \mathbf{F}_{\mathcal{S}(\mathbf{K})} \mathbf{x} \quad (1)$$

where $\mathcal{S}(\mathbf{K}) \in \Omega^{N_c \times N_s \times \frac{\Delta t}{\delta t}}$ is the sampling pattern obtained after linearly interpolating each shot in \mathbf{K} by a factor of $\frac{\Delta t}{\delta t}$. In this work we used $\delta t = 2 \mu\text{s}$. This interpolation is done to model the ADC hardware in place at the scanner. Note that at a cost of increased complexity more realistic forward models can be considered to take off-resonance effects as well as T_2^* -decay into account. However, they were not included in this work.

2.2. Reconstruction

Throughout this work, we denote the reconstruction network as $\mathcal{R}_{\theta}^{\mathbf{K}}(\mathbf{y})$, which is parameterized by θ and depends on the k-space trajectory \mathbf{K} being optimized (through the interpolated trajectory $\mathcal{S}(\mathbf{K})$) and its corresponding k-space measurements \mathbf{y} .

DCp Adjoint: As a first step, we rely on a simple DCp adjoint NUFFT to perform image reconstruction. This ensures that the given network actually results in trajectories which are data-aware and samples the center of k-space more densely as expected by traditional CS theory [1–5]. Note that in this case the reconstruction network is parameter-free and satisfies the following reconstruction formula:

$$\hat{\mathbf{x}}_{dc} = \mathcal{R}^{\mathbf{K}}(\mathbf{y}) = \mathbf{F}_{\mathcal{S}(\mathbf{K})}^{\text{H}} \mathbf{D}_{\mathcal{S}(\mathbf{K})} \mathbf{y} \quad (2)$$

where \mathbf{F}^{H} is the adjoint of NUFFT operator \mathbf{F} and $\mathbf{D}_{\mathcal{S}(\mathbf{K})}$ is the density compensators for the interpolated sampling pattern $\mathcal{S}(\mathbf{K})$ obtained iteratively using the method described in [20].

U-Net: We then used a U-net architecture [21] directly on the density-compensated adjoint image $\hat{\mathbf{x}}_{dc}$. This network can be trained to reduce aliasing and other artifacts in the final reconstructed image.

NC-PDNet: Finally, we made use of the state-of-the-art DCp unrolled network called NC-PDNet [15]. This network is known to provide improved image quality compared to U-Net and remain robust to changes in trajectories, which is crucial in our approach as the sampling pattern is learned and updated in every gradient descent step. A key difference of this network as compared to that used in BJORK [12] is that we use DCp to better condition the reconstruction problem. Moreover, in the underlying CNN-based denoiser, we use different parameters across iterations, giving improved reconstruction performance at the cost of larger memory footprint.

3. OPTIMIZATION SCHEME

We trained our networks for nearly 1.5k steps with batch size of 64 on the fastMRI training data, which was split into training and validation in a 90%-10% ratio. This enabled early stopping to prevent overfitting. The original fastMRI validation dataset was used only for later evaluation.

3.1. Costs, gradients and optimizers

We optimize for the k-space trajectory \mathbf{K} and reconstruction network $\mathcal{R}_{\theta}^{\mathbf{K}}$ as follows:

$$(\hat{\mathbf{K}}, \hat{\theta}) = \arg \min_{(\mathbf{K} \in \mathcal{Q}^{N_c}, \theta)} \mathcal{L}(\mathbf{x}, \mathcal{R}_{\theta}^{\mathbf{K}}(\mathbf{F}_{\mathcal{S}(\mathbf{K})} \mathbf{x})) \quad (3)$$

where the hardware constraints on the trajectories are specified by the feasible trajectory space \mathcal{Q}^{N_c} as described in Sec. 3.4. The loss function \mathcal{L} used in this study was inspired by [22] which is

a weighted sum of ℓ_1 , ℓ_2 and multi-scale structural similarity index (\mathcal{S}) [23]:

$$\mathcal{L}(\mathbf{x}, \widehat{\mathbf{x}}) = \alpha(1 - \mathcal{S}(\mathbf{x}, \widehat{\mathbf{x}})) + \bar{\alpha} \|\mathbf{x} - \widehat{\mathbf{x}}\|_1 + \frac{\bar{\alpha}^2}{2} \|\mathbf{x} - \widehat{\mathbf{x}}\|_2$$

with $\bar{\alpha} = 1 - \alpha$ and the value of α was tuned to 0.998 to give nearly equally balanced terms.

Optimizing the trajectory in (3) requires computing the gradient of \mathcal{L} with respect to \mathbf{K} :

$$\frac{\partial \mathcal{L}(\mathbf{x}, \widehat{\mathbf{x}})}{\partial \mathbf{K}} = \nabla \mathcal{L}(\mathbf{x}, \widehat{\mathbf{x}}) \frac{\partial \widehat{\mathbf{x}}}{\partial \mathbf{K}} = \nabla \mathcal{L}(\mathbf{x}, \widehat{\mathbf{x}}) \frac{\partial \mathcal{R}_\theta^{\mathbf{K}}(\mathbf{y})}{\partial \mathbf{K}} \quad (4)$$

In order to simplify this gradient calculation and reducing its computational complexity, we neglect the contribution of gradients from density compensators $D_{\mathcal{S}(\mathbf{K})}$. For the case of DCp Adjoint reconstruction, Eq. (4) reads:

$$\frac{\partial \mathcal{L}}{\partial \mathbf{K}} = \nabla \mathcal{L} \left(\frac{\partial \widehat{\mathbf{x}}}{\partial D_{\mathcal{S}(\mathbf{K})} \mathbf{y}} \mathbf{D}_{\mathcal{S}(\mathbf{K})} \frac{\partial (\mathbf{F}_{\mathcal{S}(\mathbf{K})} \mathbf{x})}{\partial \mathbf{K}} + \frac{\partial \mathbf{F}_{\mathcal{S}(\mathbf{K})}^H}{\partial \mathbf{K}} \right)$$

In order to obtain the gradient of NUFFT operators $\mathbf{F}_{\mathcal{S}(\mathbf{K})}$ and $\mathbf{F}_{\mathcal{S}(\mathbf{K})}^H$ with respect to the k-space trajectory \mathbf{K} , we implemented [14] in `tfkbNUFFT`. As these underlying gradients vary extremely in norm depending on the k-space region (as noted in [24]), we used the ADAM optimizer for learning the trajectories, while we relied on rectified-ADAM for optimizing the image reconstruction network $\mathcal{R}_\theta^{\mathbf{K}}$. We used a learning rate of 10^{-3} .

3.2. Multi-resolution

Following [25], we used a multi-resolution approach to optimize the k-space trajectory and allow for the algorithm to reach faster convergence. Consequently, the sampling pattern was first optimized on down-sampled trajectories. After optimizing over $N_{ds} \approx 250$ steps, the solution $\widehat{\mathbf{K}}^{n_1}$ at the resolution level n_1 was then interpolated and further used as a warm restart for the up-sampled problem $n_2 = 2n_1$. We used dyadic scaling and trained our trajectory over 6 decimation levels (i.e. $n_1 = N_c \times N_s / 2^5$ up to $n_6 = 2^5 n_1$). At any decimation level n_k , the solution $\widehat{\mathbf{K}}^{n_k}$ used in the forward model (1) was still linearly interpolated as $\mathcal{S}(\widehat{\mathbf{K}}^{n_k})$ to match the final trajectories shape and ensure that the reconstructor sees the same amount of data at every n_k .

3.3. Learning schemes

Here there are two sub-models that are being optimized: The k-space trajectory \mathbf{K} and the image reconstruction network $\mathcal{R}_\theta^{\mathbf{K}}$. These sub-models can be trained in the following ways:

Joint Learning (JL): In an ideal scenario, we would like to jointly optimize both sub-models. This idea was pushed forward pretty successfully in both PILOT and BJORK [11, 12] at the cost of producing trajectories close to their initialization. Here, as we rely on an unconstrained parameterization of \mathbf{K} this scheme may lead to poor local minimizers. The reason is that the reconstruction network $\mathcal{R}_\theta^{\mathbf{K}}$ is initially naive, and thus not trained on various types of trajectories, hence it may give poor reconstructed images associated with high losses \mathcal{L} for certain configurations of $\mathcal{S}(\mathbf{K})$. This initial bad conditioning of $\mathcal{R}_\theta^{\mathbf{K}}$ cannot be addressed by updating \mathbf{K} . Consequently, the gradient of \mathcal{L} with respect to \mathbf{K} (see (4)) is corrupted leading to trajectory updates in wrong directions.

Alternative Descent (AD): In this scenario, we first optimize the trajectory using the DCp Adjoint reconstruction method. As the

latter is parameter-free, the optimized sampling pattern $\widehat{\mathbf{K}}$ is purely data-driven and it densely samples the center of k-space, hence providing reasonable image reconstruction. Later on, we learn $\mathcal{R}_\theta^{\mathbf{K}}$. Although this method can provide improved results, it still does not truly learn the sub-models in a joint fashion.

Hybrid Learning (HL): In this third scenario, we try to retain the best of both previous schemes, by slowly moving from AD to JL. As a first step, we start off by only learning the trajectory \mathbf{K} using DCp adjoint as reconstruction. Following this, we carry alternative descent between \mathbf{K} and $\mathcal{R}_\theta^{\mathbf{K}}$ while slowly increasing the number of steps to move from noisy to more accurate gradients over time. In practice, this was done at different resolution levels as prescribed in Tab. 1.

Table 1: Hybrid learning strategy

Decimation levels	Learning Method
32 \rightarrow 4	Alternative descent for 20 steps each
4 \rightarrow 1	Alternative descent for $N_{ds}/2$ steps each
1	True joint optimization throughout

3.4. Constraints and Projection

In order to obtain realistic learned trajectories which can be run by a scanner gradient hardware, we need to impose G_{\max} and S_{\max} constraints. Under the normalized sampling domain $\Omega = [-1, 1]^2$, we obtain the feasible constraint set \mathcal{Q}^{N_c} from [8, Eq. (2)]. To enforce these constraints, we project all shots in \mathbf{K} using the projector $\Pi_{\mathcal{Q}^{N_c}}$ developed in [8, Sec. S1] after each gradient descent step (for more details, see [16]). With this in mind, given a step size η_t , each update of the trajectory along the learning process reads:

$$\mathbf{K}_{t+1} \leftarrow \Pi_{\mathcal{Q}^{N_c}} \left(\mathbf{K}_t - \eta_t \frac{\partial \mathcal{L}(\mathbf{x}, \widehat{\mathbf{x}})}{\partial \mathbf{K}} \right) \quad (5)$$

Given this projection step our trajectories strictly meet the constraints on G_{\max} and S_{\max} in contrast to earlier works [11–13] which handle these constraints on the trajectories through a penalty term in the loss \mathcal{L} . Importantly, this projection step is parameter-free. Additionally, the projector has no influence over the gradient of \mathcal{L} with respect to \mathbf{K} , allowing it to improve image quality at the reconstruction stage.

4. NUMERICAL STUDIES

We compare our results for AF = 20, and 2 imaging contrasts (T1-w, T2-w) using AD and HL as learning strategies. We did not compare with JL strategy as our experiments led to low SSIM scores (between 0.7 and 0.8 after training). Further, for comparison with an earlier baseline, we use SPARKLING trajectories generated with the learned TSD using LOUPE as obtained in [9].¹ We denote this approach by SP. The learned models were tested on 512 slices from the fastMRI validation dataset. We used the SSIM and PSNR scores on each slice as image quality metrics for evaluation.

4.1. Quantitative

We first present the trajectory and quantitative comparisons for each of the learning strategies in the form of box plots in Fig. 2. One key

¹Comparisons with PILOT [11] and BJORK [12] could not be done due to lack of time, although we thank the authors for sharing their trajectories with us. We will compare them in future work.

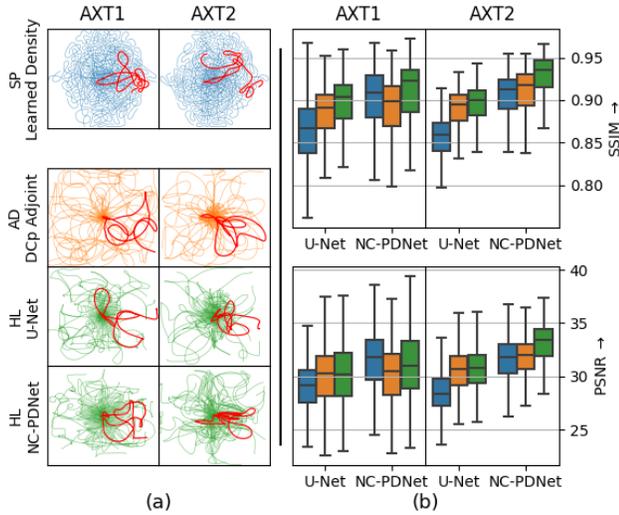


Fig. 2: (a) The optimized k-space trajectories with SPARKLING on learned density (top, blue) and by joint optimization (bottom). For joint optimization, results with AD (orange) and HL (green) methods are presented for (AX)T1/(AX)T2 contrasts. Finally, trajectories optimized with U-Net and NC-PDNet as reconstructors shown for HL method. (b) Box plots summarizing the corresponding color-matched image reconstruction results on a retrospective study using 512 slices (fastMRI validation dataset). SSIMs/PSNRs appear on top/at the bottom.

aspect of all our learned trajectories is that they are extremely curved (in contrast to [11, 12]) and maximally explore the k-space, which ideally will lead to better performance. We notice that for the AD strategy, where the trajectory was learned with parameter free DCp adjoint as reconstructor, we do sample the center of k-space densely as expected by CS theory.

We observe that NC-PDNet consistently outperforms U-Net as it is an unrolled network. Further, we see that HL learning strategy clearly outperforms all the other methods with SSIM scores in the range of 0.9-0.95. The only exception to this is PSNR values for T_1 -w contrasts, where SP presents equivalent results to NC-PDNet. However, note that the SP trajectories were obtained by learning the density, followed by obtaining trajectories with SPARKLING and finally learning the NC-PDNet network for it. HL strategy on the other hand learns jointly in a single training.

4.2. Qualitative

In order to assess image quality and the residuals across the different learning schemes (SP, AD, HL), we present a single reconstructed slice for T_1 -w images in Fig. 3 and T_2 -w images in Fig. 4. We clearly observe lesser residual structures in the HL/NC-PDNet approach, which can also be verified quantitatively with highest SSIM and PSNR scores for 0.926 and 34.241, respectively for T_1 -w images. However, it is notable to observe that SP approach gives at par results as observed previously. For T_2 -w, we see HL/NC-PDNet approach clearly outperforms with SSIM and PSNR scores of 0.95 and 34.865 respectively.

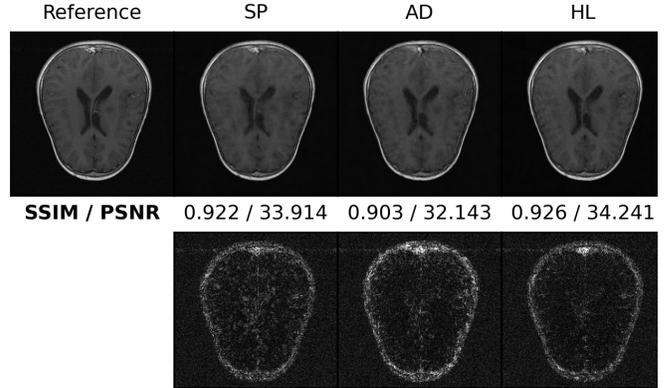


Fig. 3: Top: Reference and reconstruction results for a single slice from file_brain_AXT1.202.2020283.h5 with SP, AD and HL strategies and NC-PDNet reconstructor. **Bottom:** The residuals, scaled to match and compare across methods.

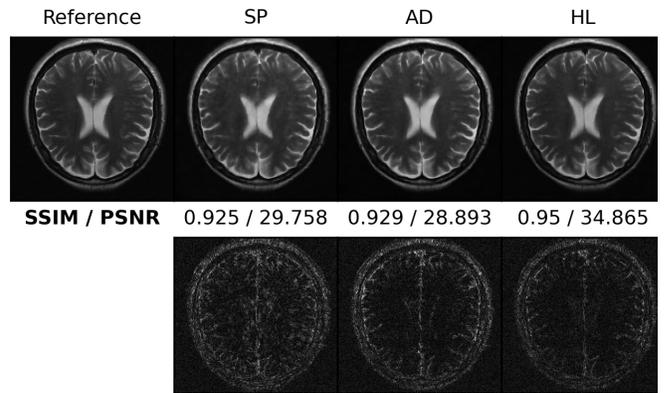


Fig. 4: Top: Reference and reconstruction results for a single slice from file_brain_AXT2.209.6001069.h5 with SP, AD and HL strategies and NC-PDNet reconstructor. **Bottom:** The residuals, scaled to match and compare across methods.

5. CONCLUSIONS

In this work, we present a hybrid framework for jointly learning the trajectory and image reconstructor through projected gradient descent, which provides trajectories comparable to those generated by SPARKLING with learned density. Through retrospective studies on the fastMRI validation dataset, we showed that this novel learning scheme works across multiple resolutions and leads to superior performance of the trajectories and improved image quality overall.

Future prospects of this work include prospective implementations through modifications of T_1 and T_2 -w imaging sequences. This could lead to additional constraints on the trajectory arising from the MR sequencing or contrasts.

Acknowledgments

This work was granted access to the HPC resources of IDRIS under the allocation 2021-AD011011153 made by GENCI. Chaithya G R was supported by the CEA NUMERICS program, which has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 800945.

6. REFERENCES

- [1] G. Puy *et al.*, “On variable density compressive sampling,” vol. 18, no. 10, pp. 595–598, 2011.
- [2] N. Chauffert *et al.*, “Variable density compressed sensing in mri. theoretical vs heuristic sampling strategies,” in *2013 IEEE 10th International Symposium on Biomedical Imaging*. IEEE, 2013, pp. 298–301.
- [3] —, “Variable density sampling with continuous trajectories. Application to MRI,” vol. 7, no. 4, pp. 1962–1992, Nov. 2014.
- [4] B. Adcock, A. C. Hansen, C. Poon, B. Roman *et al.*, “Breaking the coherence barrier: asymptotic incoherence and asymptotic sparsity in compressed sensing,” *arXiv preprint arXiv:1302.0561*, 2013.
- [5] C. Boyer *et al.*, “Compressed sensing with structured sparsity and structured acquisition,” vol. 46, no. 2, pp. 312–350, 2019.
- [6] C. Lazarus *et al.*, “SPARKLING: variable-density k-space filling curves for accelerated T2*-weighted MRI,” vol. 81, no. 6, pp. 3643–3661, 2019.
- [7] —, “3D variable-density SPARKLING trajectories for high-resolution T2*-weighted magnetic resonance imaging,” 2020.
- [8] G. R. Chaithya *et al.*, “Optimizing full 3D SPARKLING trajectories for high-resolution T2*-weighted Magnetic Resonance imaging,” *under review at IEEE Transactions on Medical Imaging*, May 2021. [Online]. Available: https://hal.inria.fr/hal-03090471v2/file/2020_Full3D_SPARKLING.pdf
- [9] —, “Learning the sampling density in 2D SPARKLING MRI acquisition for optimized image reconstruction,” May 2021, working paper or preprint. [Online]. Available: <https://hal.inria.fr/hal-03158831>
- [10] C. D. Bahadir *et al.*, “Deep-learning-based optimization of the under-sampling pattern in MRI,” vol. 6, pp. 1139–1152, 2020.
- [11] T. Weiss *et al.*, “PILOT: Physics-informed learned optimal trajectories for accelerated MRI,” *arXiv:1909.05773v4*, Aug. 2020.
- [12] G. Wang *et al.*, “B-spline parameterized joint optimization of reconstruction and k-space trajectories (BJORK) for accelerated 2d MRI,” *arXiv preprint arXiv:2101.11369*, 2021.
- [13] S. Vedula *et al.*, “3D FLAT: Feasible Learned Acquisition Trajectories for Accelerated MRI,” in *Machine Learning for Medical Image Reconstruction: 3rd Intern. WS MLMIR 2020, Held in Conjunction with MICCAI 2020*. Lima, Peru: Springer Nature, Oct. 2020, p. 3.
- [14] W. Guanhua, C. N. Douglas, and J. A. Fessler, “Efficient NUFFT Backpropagation for Stochastic Sampling Optimization in MRI,” in *29th Proceedings of the ISMRM society*, no. 0913, virtual, May 2021.
- [15] Z. Ramzi *et al.*, “NC-PDNet: a Density-Compensated Unrolled Network for 2D and 3D non-Cartesian MRI Reconstruction,” Sep. 2021, this work is an extended version of the work presented at the 2021 ISBI conference. [Online]. Available: <https://hal.inria.fr/hal-03188997>
- [16] N. Chauffert *et al.*, “A projection algorithm for gradient waveforms design in MRI,” vol. 35, no. 9, pp. 2026–2039, Sep. 2016.
- [17] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [18] J. Zbontar *et al.*, “fastMRI: An open dataset and benchmarks for accelerated MRI,” *arXiv:1811.08839*, 2018.
- [19] D. L. Parker *et al.*, “Phase reconstruction from multiple coil data using a virtual reference coil,” *Magnetic Resonance in Medicine*, vol. 72, no. 2, pp. 563–569, 2014.
- [20] J. G. Pipe and P. Menon, “Sampling density compensation in MRI: Rationale and an iterative numerical solution,” vol. 41, no. 1, pp. 179–186, 1999.
- [21] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” 2015.
- [22] N. Pezzotti *et al.*, “An adaptive intelligence algorithm for undersampled knee mri reconstruction,” *IEEE Access*, vol. 8, pp. 204 825–204 838, 2020.
- [23] Z. Wang, E. Simoncelli, and A. Bovik, “Multiscale structural similarity for image quality assessment,” in *The Thirty-Seventh Asilomar Conference on Signals, Systems Computers, 2003*, vol. 2, 2003, pp. 1398–1402 Vol.2.
- [24] D. Gournay *et al.*, “Spurious minimizers in non uniform Fourier sampling optimization,” Jun. 2021, working paper or preprint. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03212145>
- [25] L. Lebrat *et al.*, “Optimal Transport Approximation of 2-Dimensional Measures,” vol. 12, no. 2, pp. 762–787, Jan. 2019.