



HAL
open science

NC-PDNet: a Density-Compensated Unrolled Network for 2D and 3D non-Cartesian MRI Reconstruction

Zaccharie Ramzi, Chaithya Giliyar Radhakrishna, Jean-Luc Starck, Philippe Ciuciu

► **To cite this version:**

Zaccharie Ramzi, Chaithya Giliyar Radhakrishna, Jean-Luc Starck, Philippe Ciuciu. NC-PDNet: a Density-Compensated Unrolled Network for 2D and 3D non-Cartesian MRI Reconstruction. IEEE Transactions on Medical Imaging, 2022, 10.1109/TMI.2022.3144619 . hal-03188997v3

HAL Id: hal-03188997

<https://inria.hal.science/hal-03188997v3>

Submitted on 31 Jan 2022

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.

NC-PDNet: a Density-Compensated Unrolled Network for 2D and 3D non-Cartesian MRI Reconstruction

Zaccharie Ramzi, Chaithya G R, Jean-Luc Starck and Philippe Ciuciu *Senior Member, IEEE*

Abstract—Deep Learning has become a very promising avenue for magnetic resonance image (MRI) reconstruction. In this work, we explore the potential of unrolled networks for non-Cartesian acquisition settings. We design the *NC-PDNet* (Non-Cartesian Primal Dual Network), the first density-compensated (DCp) unrolled neural network, and validate the need for its key components via an ablation study. Moreover, we conduct some generalizability experiments to test this network in out-of-distribution settings, for example training on knee data and validating on brain data. The results show that *NC-PDNet* outperforms baseline (U-Net, Deep image prior) models both visually and quantitatively in all settings. In particular, in the 2D multi-coil acquisition scenario, the *NC-PDNet* provides up to a 1.2dB improvement in peak signal-to-noise ratio (PSNR) over baseline networks, while also allowing a gain of at least 1dB in PSNR in generalization settings. We provide the open-source implementation of *NC-PDNet*, and in particular the Non-uniform Fourier Transform in TensorFlow, tested on 2D multi-coil and 3D single-coil k-space data.¹

Index Terms—MRI, Image Reconstruction, Deep Learning, Non-Cartesian, Unrolled Networks.

I. INTRODUCTION

Magnetic Resonance Imaging (MRI) is the reference imaging technique used to probe soft tissues in the human body non-invasively. The data acquisition process in MRI is inherently slow due to the sequential measurements collection in the Fourier domain, also called k-space. This slow acquisition causes many issues such as motion-corrupted scans, limited image resolution in a scan time compatible with clinical routine, non-applicability to certain patients, and limited patient throughput. To reduce the scan time as a means to offset the limiting factors, the main solution proposed in the MRI literature was to diminish the number of measurements in the k-space. It was exploited for parallel imaging (PI) [2], [3] as well as for Compressed Sensing (CS) [4]. In the latter framework, a hand-crafted decomposition basis is used to express the sparsity of the image to be reconstructed in a transform domain [5], [6]. From this sparsity prior, an optimization problem, with a solution close to the image of interest, can be constructed.

Zaccharie Ramzi, Chaithya G R and P. Ciuciu are with CEA, NeuroSpin, F-91191 Gif-sur-Yvette, cedex, France and Inria, Parietal, Université Paris-Saclay, F-91120 Palaiseau, France (e-mail: zaccharie.ramzi@inria.fr), whereas J.-L. Starck is affiliated with AIM, CEA, CNRS, Université Paris-Saclay, Université Paris Diderot, Sorbonne Paris Cité F-91191 Gif-sur-Yvette, France.

¹This work is an extended version of the work presented at the 2021 ISBI conference [1].

This optimization problem stems from casting the reconstruction as an idealized inverse problem (i.e. ignoring magnetic gradient inaccuracies, B_0 field inhomogeneities, T_2 decay and off-resonance effects). It can be written in its most general form as:

$$\mathbf{y}_\ell = \mathcal{F}_\Omega \mathbf{S}_\ell \mathbf{x} + \mathbf{n}_\ell \quad \forall \ell = 1, \dots, L \quad (1)$$

with L the number of coils used for the acquisition, \mathcal{F}_Ω the measurements operator, a 2D or 3D Non-uniform Discrete Fourier Transform (NDFT) [7] with locations defined by Ω (if Ω is on a grid, $\mathcal{F}_\Omega = \Omega \mathbf{F}$ can be simplified to a regular Fourier Transform \mathbf{F} followed by a masking operator parameterized by Ω). Diagonal matrices $(\mathbf{S}_\ell)_{\ell=1}^L$ define the sensitivity maps (in the case of a single coil acquisition, $L = 1$ and \mathbf{S}_1 is just the identity matrix), $\mathbf{y}_\ell \in \mathbb{C}^m$ are the k-space measurements associated with each coil, \mathbf{n}_ℓ is an additive white Gaussian noise (AWGN), and $\mathbf{x} \in \mathbb{C}^n$ the target image to reconstruct. To make scan times shorter, PI and CS are of course associated with the same requirement, namely $m \ll n$, even though they both differ in the way under-sampling is implemented (regular for PI vs. incoherent for CS).

The reconstructed image $\hat{\mathbf{x}}$ is then recovered by solving the following optimization problem:

$$\arg \min_{\mathbf{x} \in \mathbb{C}^n} \sum_{\ell=1}^L \frac{1}{2} \|\mathbf{y}_\ell - \mathcal{F}_\Omega \mathbf{S}_\ell \mathbf{x}\|_2^2 + \mathcal{R}(\mathbf{x}) \quad (2)$$

where \mathcal{R} is a regularization term that serves as a prior of the solution. Usually, \mathcal{R} is hand-crafted and can be the ℓ_1 -norm of the decomposition of \mathbf{x} in some adequate basis like the wavelets, or the gradient of the image (TV regularization).

A prototypical version of an iterative algorithm that can solve the optimization problem (2) is the following proximal gradient descent:

$$\begin{aligned} \mathbf{x}_{n+1} &= \mathbf{x}_n - \epsilon_n \mathcal{A}^H (\mathcal{A} \mathbf{x}_n - \mathbf{y}) \\ \mathbf{x}_{n+1} &= \text{prox}_{\epsilon_n \mathcal{R}} (\mathbf{x}_{n+1}) \end{aligned} \quad (3)$$

where $\mathcal{A} = (\mathbf{I}_L \otimes \mathcal{F}_\Omega) \mathcal{S}$, $\mathcal{S} = [\mathbf{S}_1^H, \dots, \mathbf{S}_L^H]^T$, $\mathbf{y} = [\mathbf{y}_1^H, \dots, \mathbf{y}_L^H]^T$, ϵ_n is the gradient step size at the n -th iteration, and H is the Hermitian (i.e. transpose conjugate) operator. The following issues can be spotted using this method: (i) iterations have computationally-intensive operators, (ii) the prior \mathcal{R} needs to be manually crafted and ideally its proximal operator has to be closed form. While dictionary learning methods [8] can alleviate the second issue, they do so at an increased computation cost but still with relatively simple (i.e. linear) learned priors.

To circumvent these problems, Deep Learning (DL) methods have become of great interest for MRI reconstruction. Indeed, the unrolled framework [9] allows for the use of

less iterations for image reconstruction while enabling the neural network to learn more complex priors on MR images from the data itself. Hence, it tackles both problems (i)-(ii) observed in CS methods at the same time. Moreover, the current deep learning frameworks offer an off-the-shelf efficient implementation on GPU and TPU which also facilitates gains in terms of reconstruction speed. While the first deep learning methods applied to MRI reconstruction were not necessarily using the unrolled framework [10], the recent state-of-the-art methods all rely on this framework [11]–[14]. Most of these approaches are employed to reconstruct images from k-space data collected on a Cartesian grid, and therefore cannot be tested in the more general non-Cartesian acquisition setting, when the measurements are collected off the grid.

Non-Cartesian MRI acquisitions are of interest for multiple purposes. For example, the radial acquisition that happens with projection acquisition can be used for shorter echo times, reduced sensitivity to motion, and improved temporal resolution in some applications (e.g. angiography, dynamic MRI) [15]–[17]. The same statement holds for spiral imaging that remains insensitive to in-plane flow-related artifacts [16], [18]. As these setups are used in clinical routine, it is crucial to push the development of DL reconstruction methods for their application to these types of data.

Moreover, the interest in going from 2D to 3D imaging can be justified by a higher signal-to-noise ratio (SNR) requested for targeting higher resolution in clinical use [19]. Although multiple approaches exist to efficiently sample the 3D k-space [17], [20]–[22], we will focus on full 3D radial under-sampling as the stacking strategies (e.g. stack of stars) can be tackled using a 2D (i.e. slicewise) reconstruction method. However, due to hardware constraints, we focus on single-coil 3D imaging rather than on multi-coil 3D imaging, and leave this extension for future work.

In this paper, we contribute to the domain of non-Cartesian MRI reconstruction by:

- introducing the *NC-PDNet* (i.e. Non-Cartesian PDNet), the first density-compensated unrolled neural network for non-Cartesian k-space data;
- implementing the first unrolled networks for non-Cartesian multi-coil MRI 2D and single-coil MRI 3D data;
- open sourcing our implementation for all the networks, data processing, training, and evaluation² in TensorFlow [23], in particular a version of the Non-Uniform Fast Fourier Transform (NUFFT) in the TensorFlow framework, usable for multi-coil and 3D data (with the corresponding density compensation code)³;
- performing some out-of-distribution performance tests for the trained networks in multiple acquisition scenarios and comparing them against the state of the art.

II. RELATED WORKS

a) Correcting k-space for non-Cartesian reconstruction: This work strongly builds on the seminal paper by *Pipe and*

Menon [24] introducing density compensation (DCp⁴) for non-Cartesian k-space data. In this work, the authors provide a solution for computing the DCp which avoids some requirements that may be difficult to meet in practice, such as the knowledge of the k-space readouts or the constraint associated with the Nyquist criterion. However, this work does not use a learning based reconstruction strategy to complement the application of the DCp.

Some other works have explored other strategies to correct the k-space for non-Cartesian reconstruction, using a pre-conditioning of the k-space in iterative CS procedures in order to accelerate the convergence (not necessarily to obtain better image quality for an unlimited computational budget) [25]. To improve the latter approach, the authors in [26] introduced a diagonal pre-conditioner optimized for the Mean Squared Error (MSE) that did not need an inner loop. However, these works use classical iterative algorithms which have been shown to be outperformed by end-to-end learned unrolled approaches in various settings [11], [27]–[29].

b) Classical non-Cartesian reconstruction: Classical reconstruction algorithms like SPIRiT and GRAPPA [3] have been adapted to the non-Cartesian setting and a summary of these adaptations is available in [30].

Another typical setting where non-Cartesian reconstruction is needed is dynamic MRI reconstruction. In this context, many works have used model-based optimization to reconstruct MR images from non-Cartesian acquisitions. The STORM method [31] assumes that the different frames of the dynamic MRI acquisition evolve in the same low-dimensional manifold. The core idea is the following: As the underlying organ does not evolve too much in time, navigators (i.e. a calibration signal) can be used to recognize similar frames and build a Laplacian matrix from the latter in order to obtain a manifold regularizer. This idea is not applicable to the acquisition scenarios we cover in this paper as it strongly relies on the temporal nature of the signal to build a fitting regularization. In contrast, we consider here non-Cartesian reconstruction as a standalone task. Works derived from STORM or similar to it [32]–[34] also make use of the temporal nature of dynamic MRI to introduce a manually crafted regularization.

c) Unrolled neural networks: The idea of unrolling an optimization algorithm was originally introduced for the problem of sparse coding [9]. It has then been applied to many different inverse problems and in particular in MRI [11], [12], [29], [35]. While an unrolled neural network architecture can be created for all the different iterative algorithms solving Eq. (2) (e.g. ADMM [36], Neumann Series [37], Dictionary Learning [29] or Alternating Minimization [28]), in this work we focus only on one of them, the XPDNet framework [38], as it is the highest ranked framework with an open-source implementation in the fastMRI 2020 reconstruction challenge [14], giving the name to our network: the *NC-PDNet*. The aim of this paper is not actually to compare different strategies for the unrolled architecture but instead to demonstrate the improvement inherent to the density compensation mechanism

²<https://github.com/zaccharieramzi/fastmri-reproducible-benchmark>

³<https://github.com/zaccharieramzi/tfkbnuft>

⁴The acronym DC is usually used to denote Data Consistency in the MRI reconstruction field.

for a given unrolled architecture.

d) Neural networks for non-Cartesian reconstruction: The Nonuniform Variational Network [29] was the first unrolled network to be designed for non-Cartesian MRI reconstruction. This network unrolls a proximal gradient descent algorithm and operates on 192×192 single-coil images with a variable density acquisition scheme. This network, however, does not include a DCp scheme, which is most likely due to the fact that the trajectories studied in this work do not present a variable density as recommended in the CS literature [39]–[41]. Additionally, this network was only applied to single-coil 2D data whose phase was simulated. Moreover, there is no open-source implementation for this work to date. In contrast, our network is applied to both single and multi-coil 2D raw k-space data as well as to 3D data without phase information.

A perceptual complex neural network was designed in [42] to tackle 2D + time real-time cine MRI in cardiac imaging. However, this architecture is not an unrolled network and was applied to the gridding-reconstructed MR image. The authors of [42] specify that they did not consider data consistency-based methods due to the increased computation cost of the NUFFT. They actually pointed out that efficient tools like `torchkbnufft` [43] would be helpful in this regard. Here, we show that using a DCp step, a modeling choice allowed by an efficient implementation of the NUFFT, is actually critical to obtain improved results compared to baseline neural networks.

Similarly to our work, in [44] an unrolled network is used to reconstruct non-Cartesian MR angiography data. While they did use a density compensation scheme, they did not perform an ablation study, and did not use modern techniques for MRI unrolled networks such as buffers or sensitivity maps refinement. Another aspect is that the way the algorithm is unrolled in [44] embeds linearly the data consistency, formally its i -th iteration is of the form $\mathbf{x}_{i+1} = f_{\theta_i}(\mathbf{x}_i - \epsilon_i \mathcal{A}^H(\mathcal{A}\mathbf{x}_i - \mathbf{y}))$. In contrast, similarly to [45], the learned Primal Dual algorithm [27] used in this work allows us to mix the prior information and the data consistency in a nonlinear way. Its i -th iteration reads as follows: $\mathbf{x}_{i+1} = f_{\theta_i}([\mathbf{x}_i; \mathcal{A}^H(\mathcal{A}\mathbf{x}_i - \mathbf{y})])$. This second form could enable more easily to compensate the density on its own, because the model can act separately on the data consistency component $\mathcal{A}^H(\mathcal{A}\mathbf{x}_i - \mathbf{y})$. We show via our ablation studies that we still need the density compensation weights even with this modern unrolling technique. Furthermore, their code is not available online neither for the network, nor for the NUFFT.

Another direction of research for non-Cartesian MRI reconstruction is the use of untrained networks with the Deep Image Prior (DIP) [46] framework. In [47], the authors designed a DIP approach tailored for dynamic MRI using a manifold to sample from for the random input. While this specificity is not applicable to our setting, it is a strong contender in its time independent (i.e. static) version, especially for the out-of-distribution performance evaluations. For this reason, we will perform a comparative analysis to the DIP framework.

e) Neural networks for 3D reconstruction: In [48], the authors introduced a complex unrolled network for multi-coil 3D + time reconstruction with two unrolled iterations for

Cartesian acquisitions. However, due to the memory requirements of the NUFFT, it is not clear how to adapt this solution to the non-Cartesian case. Multi-coil 3D MRI reconstruction is also tackled in [49], where the authors design a memory-efficient algorithm to train unrolled networks, however again this approach was restricted to Cartesian acquisitions.

III. MODEL

A. Unrolled network

In this work, we unroll the proximal gradient descent (3) using the alternate scheme proposed in [27] (that is the primal only version of the network).

Algorithm 1: *NC-PDNet*: Density compensated Primal Dual unrolled neural network over N_C iterations.

Data: \mathbf{y} the k-space data, Ω the non-Cartesian trajectories, \mathbf{d} the pre-computed DCp weights, \mathcal{S} the coarse estimates of the sensitivity maps

Result: \mathbf{x} , the reconstructed magnitude MR image
 $\mathcal{S} = g_{\theta_r}(\mathcal{S})$; // Sensitivity maps refinement
 Update \mathcal{A} and \mathcal{A}^H ;

$\mathbf{x} = \mathcal{A}^H \mathbf{y}$;
 $\mathbf{x}_b = [\mathbf{x}, \mathbf{x}, \mathbf{x}, \mathbf{x}, \mathbf{x}]$; // Buffer creation, in practice concatenation along the channel dimension

for $i \leftarrow 1$ **to** N_C **do**
 $\mathbf{y}_{res} = \mathcal{A} \mathbf{x}_b[0] - \mathbf{y}$; // Data consistency
 $\mathbf{x}_{dc} = \mathcal{A}^H \mathbf{d} \mathbf{y}_{res}$; // Density compensation
 $\mathbf{x}_b = \mathbf{x}_b + f_{\theta_i}([\mathbf{x}_b, \mathbf{x}_{dc}])$; // Proximity operator learning and nonlinear acceleration scheme

end
 $\mathbf{x} = |\mathbf{x}_b[0]|$ // Magnitude computation

We describe the algorithm for the forward pass of the *NC-PDNet* in Alg. 1, in the general multi-coil setting, and an illustration of its architecture is depicted in Fig. 1. The only adjustment to be made for the single-coil case, is that sensitivity maps \mathcal{S} are no longer involved.

In Alg. 1, each network f_{θ_i} (called image net in Fig. 1) is a convolutional neural network consisting of 3 convolutions applied successively to the input, with a ReLU activation function applied after each of the first two convolutions (as was done in [27]). We specify here that the number of unrolled iterations N_C is fixed, that each network f_{θ_i} has its own set of weights irrespective of the other (f_{θ_j}) $_{j \neq i}$, and that the whole network is trained end-to-end. This makes a significant difference compared to the plug-and-play framework [50] where the number of iterations is not fixed, each sub-network has the same weights and the sub-network is learned offline.

a) Iterations Buffer: Following the idea of [27] (called “memory” in the original paper [51]), we carry a buffer \mathbf{x}_b between the different unrolled iterations rather than a single estimate. This buffer of size N_P allows us to learn a nonlinear acceleration scheme as explained in [51].

b) *Sensitivity maps refinement*: g_{θ_r} is a U-net [52] applied uniformly to all the sensitivity maps for their improved estimation (i.e. refinement) [53]. The coarse estimates of the sensitivity maps \mathcal{S} are computed using the density compensated NUFFT adjoint of the low frequencies of \mathbf{y} (i.e. typically the central part of k-space corresponding to 20 % of data), rather than using the ESPIRiT method [54]. Indeed ESPIRiT is computationally intensive and any approximation could be corrected by the network. The U-net we use is rather small with 3 scales and 4 base filters. We also use a Leaky ReLU activation as in [53]. This results in the sensitivity maps refinement module being only 7k parameters big. We highlight that this module is trained in an end-to-end fashion as recommended by [53].

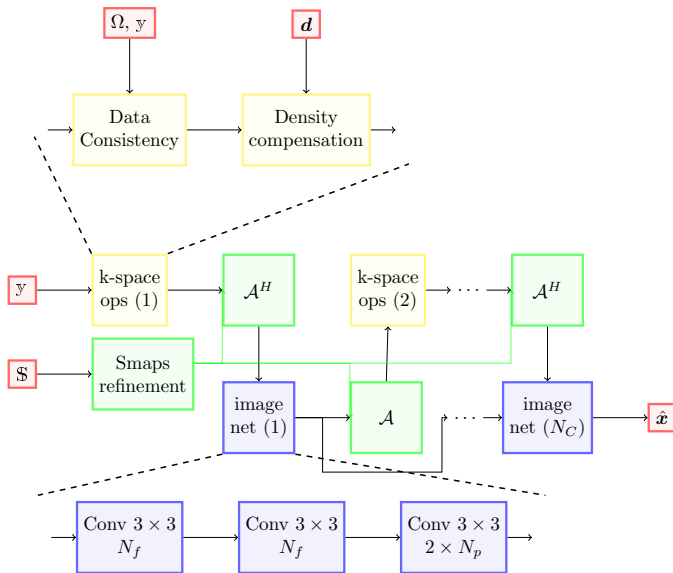


Fig. 1. Illustration of the N_C -PDNet for multi-coil MRI reconstruction. N_f is the number of filters used for the convolutions, N_p is the buffer size, and N_C the number of unrolled iterations.

c) *Complex values*: We use real-valued convolutions in the network rather than custom-made complex convolutions [55]. Indeed, real-valued convolutions turn out to have more capacity than complex-valued ones when applied to complex tensors handled by concatenating the real and imaginary parts. The concatenation is done on the channel dimension.

d) *3D*: For the 3D architecture, we replace the 2D convolutions by 3D convolutions with a kernel size of $3 \times 3 \times 3$. Because the activations of the convolutions are much larger in this case, a smaller network is designed to fit into memory when backpropagating. Therefore, we reduce the number of filters for each convolution to $N_f = 16$ (instead of $N_f = 32$ in 2D) and we only consider $N_C = 6$ (instead of $N_C = 10$ in 2D) unrolled iterations, see details in Section V-A.

B. Non-Uniform Fourier Transform

The NDFT is the generalization of the discrete Fourier Transform to positions in the Fourier space that are off the Cartesian grid and not necessarily equispaced. An approximate algorithm to have an efficient computation of the NDFT

was introduced in [7], [56]. This algorithm uses an over-sampled grid and an optimal interpolation to perform the NDFT efficiently at the cost of an approximation. We refer to this algorithm as the Non-Uniform Fast Fourier Transform (NUFFT) and highlight that unlike the Fast Fourier Transform (FFT), it is not an exact algorithm.

While the NUFFT is a more efficient algorithm than the direct application of the definition of the NDFT, it is still a very computationally demanding algorithm compared to the FFT for the same image dimensions. An alternative to the NDFT could be to grid the k-space measurements and simply use the FFT to compute the data consistency. In practice, this would mean using $\mathbf{y}_{grid} = grid(\mathbf{y}, \Omega)$ (where $grid$ is a linear gridding operation) instead of \mathbf{y} and \mathcal{F}_Ω would simply be the FFT.

a) *NDFT vs. NUFFT*: We draw the reader attention to the potential confusion between the 2 acronyms. The NDFT is the original transform, while the NUFFT is the approximate fast algorithm to compute the transform.

C. Data Consistency

The data consistency is the step in the unrolled network allowing us to inject the initial measurements \mathbf{y} by comparing them to the current estimate's $\mathbf{x}_b[0]$ measurements. The formula we have chosen for data consistency stems from the AWGN model, by taking the gradient of the ℓ_2 -norm in (2):

$$\mathbf{x}_{dc} = \mathcal{A}^H \mathbf{d} (\mathcal{A} \mathbf{x}_b[0] - \mathbf{y}). \quad (4)$$

D. Density Compensation

Unlike the Cartesian case, the adjoint operator of the NDFT is not always its inverse operator. Worse, in most cases, the NDFT does not admit an inverse operator. The application of the adjoint Fourier operator \mathcal{F}_Ω^H to the single coil k-space data (or its multi-coil extension \mathcal{A}^H to the multi-coil data) can therefore be very far from the solution to Eq. (1).

To circumvent this, DCp has been introduced [24]. Indeed, the main problem with the classical non-Cartesian trajectories like radial, spiral or any variable density sampler [39], is that they densely sample the center of k-space. Therefore, when computing the adjoint operator (\mathcal{F}_Ω^H or \mathcal{A}^H , respectively) a lot of weight is assigned to the densely sampled region at the center of the k-space, resulting in an image with abnormally large values. DCp is just the action of using factors that weigh the different sample locations so that they all play an even role during the application of the adjoint. DCp is particularly needed for deep learning approaches because the values entering convolutional layers need to have normalized (or close to normalized) values in order to avoid numerical issues.

In practice, for both the radial and spiral trajectories, we obtain the DCp factors \mathbf{d} by applying the adjoint and forward operators⁵ iteratively for N iterations, starting from $\mathbf{d}_0 = 1$:

$$\mathbf{d}_{n+1} = \frac{\mathbf{d}_n}{\mathcal{F}_\Omega \mathcal{F}_\Omega^H \mathbf{d}_n} \quad (5)$$

where the division is here pointwise. The final weights are \mathbf{d}_N . In practice, we took $N = 10$.

⁵In practice we use the interpolation operator of the NUFFT rather than the whole operator.

IV. DATA

We used the *NC-PDNet* on different data acquisition scenarios. In all of them, the data was retrospectively under-sampled using fixed non-Cartesian trajectories. Specifically in this work we used multi-shot spiral and radial trajectories for 2D imaging, while we restricted our numerical studies to the use of full 3D radial spokes for 3D imaging. A schematic view (i.e. with a much higher undersampling of the trajectories) is available in Fig. 2. The full 3D radial under-sampling pattern is illustrated in Supplementary Material in Fig. S1. In 2D, most of the numerical experiments were conducted with an acceleration factor (AF) of 4, up to an exception where $AF = 8$, compared to the full Cartesian acquisition. This factor is defined as follows for 2D imaging: $AF = \frac{N}{N_s}$ where $n = N \times N$ is the image dimension (N is the base resolution) and N_s is the number of shots involved in the under-sampling pattern. This basically means that $m = n/4$ as we performed only retrospective studies.⁶ For 3D imaging, $AF = \frac{N \times N_z}{N_s}$ where $n = N^2 \times N_z$ is the volume dimension and N_z the number of slices. In this setup, we also chose the value of $AF = 4$.




Fig. 2. Schematic representation of the two multi-shot non-Cartesian (radial and spiral) readouts considered here for 2D imaging. In our setting, we used $N_s = 100$ shots, each of them consisting of 640 samples giving a total of $m = 64,000$ k-space measurements. Here only 9 shots are presented.

A. fastMRI

For the 2D single-coil and multi-coil cases, we used the raw knee data from the fastMRI dataset [57]. The acquisition was performed using a Cartesian 2D Turbo Spin Echo (TSE) sequence and the data was collected over a 15-channel phased array coil. The imaging contrasts were proton-density weighting, half with fat suppression, half without. The sequence parameters were defined as follows: Echo train length 4, matrix size 320×320 , in-plane resolution $0.5\text{mm} \times 0.5\text{mm}$, slice thickness 3mm, no gap between slices. In total, there were 973 volumes (34,742 slices) for the training subset and 199 volumes (7135 slices) for the validation subset. In this study, we chose to jointly consider the two weighting contrasts (Proton Density and Proton Density with Fat Suppression) available in the fastMRI dataset for training. The single-coil data was emulated from the multi-coil data, meaning that it was not actually acquired with a single receiver coil. Instead, the single-coil data is the result of a linear combination of multi-coil data in

⁶In CS prospective acquisitions, oversampling can be applied along each shot to fulfill Nyquist criteria without increasing the scan time. This would make the under-sampling factor n/m lower compared to AF.

such a way that the reconstructions using root sum-of-squares match [57]. To obtain the non-Cartesian measurements \mathbf{y} , we simply used the inverse Fourier transform (denoted as \mathbf{F}^T) of the full Cartesian k-space data $\mathbf{y}_{or,\ell}$ for each coil ℓ separately before applying the NUFFT according to the non-Cartesian under-sampling pattern Ω as follows:

$$\forall \ell = 1, \dots, L \quad \mathbf{y}_\ell = \mathcal{F}_\Omega \mathbf{F}^T \mathbf{y}_{or,\ell} \quad (6)$$

We did not use the fastMRI data to perform the 3D experiments for the following reason. The fastMRI dataset was collected using a 2D multislice acquisition that uses anisotropic resolution to maintain a good SNR (in-plane resolution $0.5\text{mm} \times 0.5\text{mm}$, slice thickness 3mm). Consequently, the number of slices is very limited (roughly 30 to 40 for each acquisition) and the 3D networks will therefore not be able to fully take advantage of the third dimension.

B. OASIS

The Open Access Series of Imaging Studies (OASIS) brain database [58] includes MRI scans of 1068 participants, yielding 2168 MR sessions (3400 volumes). Of these 2168, we selected 2164 sessions which feature T1-weighted sequences. 1878 of these were acquired on a 3T scanner, 236 at 1.5T, and the remaining are undisclosed (50). We decided to use the OASIS dataset with data recast to matrix size of $176 \times 256 \times 256$, using zero-padding. While the data from the fastMRI dataset is the raw k-space measurements, the data available in the OASIS dataset is the DICOM magnitude-only image. Therefore the NDFT was applied to the magnitude volume and we did not simulate the phase, a very challenging task according to [59]. The 3D fully-sampled dataset provided in [60] was considered as a potential alternative as it contains raw data collected at isotropic resolution. Unfortunately it is not sizeable enough for deep learning tasks as specified in [57]. Consequently, this basically means that the learned model from OASIS database cannot likely be used as is in a clinical setting, and would need to be further validated, even fine-tuned. The dataset was separated in training and validation sets. The split was participant-based, that is a participant cannot have a scan in both sets. The split was of 90% for the training set and 10% for the validation set.

V. RESULTS

A. Experimental setup

a) *Metrics*: The metrics used for evaluating the networks are peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) [61]. They are computed over whole volumes of the validation set, that is the dynamic range (used for the computation of PSNR and SSIM) is computed over the whole volume. The parameters and definitions used are the same as in the fastMRI paper [57].

In order for the reader to get a better sense of the distribution of the quantitative image quality, we also provide detailed box-plots for all the quantitative results. For the sake of clarity, we chose to separate the different contrasts which present different quantitative metrics ranges in the box plots.

b) *Comparison*: To illustrate the need for all the key components of the *NC-PDNet*, we carried out an ablation study against other neural networks where we removed some aspects of the *NC-PDNet*.

- An unrolled network without DCp but only a normalization mechanism (without any normalization the network does not train due to very high values), to show the need for the DCp.
- An unrolled network without NUFFT and instead a gridded version of the k-space (the operator is only the much faster FFT), to show the need for a better approximation of the NDFT.
- A U-net [52] applied directly to the density-compensated adjoint of the k-space measurements, to show the need for the unrolling framework (3D convolutions and pooling operations are used in the 3D setting).
- The density-compensated adjoint of the k-space measurements as very naive baseline.

For the 2D settings, the unrolled networks have $N_C = 10$ unrolled iterations, use a buffer size $N_P = 5$, and a number of convolution filters $N_f = 32$. For the 3D setting, we needed to reduce those numbers for the network to fit on a single GPU, and used $N_C = 6$ unrolled iterations, a buffer size $N_P = 2$, and $N_f = 16$. The U-net was trained residually and had a base number of filters of 16.

We carried out the training and evaluation of the variants of the unrolled networks only in the 2D single-coil setting in order to save some computation time since we observed that the unrolled variants were not obtaining good performances.

We also compared to DIP, an approach that was shown to reach state-of-the-art results in dynamic MRI reconstruction [47], and that will help grasp better the generalization capacities of the networks as it is an untrained method. We specify that in order to keep the reconstruction time reasonable, we used the method advocated by [62]. This means that for a given contrast, AF, trajectory and coil number, we reuse previously trained weights in order to initialize the weights of the DIP. Similarly to prior observations in [62], we found that this strategy allows us to reduce the number of epochs by a factor of 10. For the multi-coil data, we also used the method advocated by [62] to generate the coil images with the same network, and aggregate them using root-sum-of-squares (RSS).

c) *Generalizability*: We also studied the generalizability of the trained networks to other settings, first to different sampling trajectories. In practice, we evaluated the networks trained on the spiral trajectory in the 2D multi-coil setting on the radial under-sampling pattern and vice-versa. We term this the *reverse setting*.

Second, we studied the generalizability of the networks to different organs and contrasts simultaneously. We took advantage of the 2D multi-coil brain dataset available in fastMRI to carry out this analysis. This dataset features brain data acquired with four different weighting contrasts (T1, T2, FLAIR and T1 after Gadolinium injection), collected on the same scanners, which allows us to test for generalizability solely on organ/contrast. Due to this large number of contrasts and the diversity in the number of coils for the brain data, we

TABLE I

MEAN PSNR / SSIM ON THE VALIDATION VOLUMES OF THE DIFFERENT APPROACHES AVERAGED OVER BOTH CONTRASTS (KNEE FASTMRI) IN THE SINGLE-COIL SETTING. THE BEST RESULTS ARE IN BOLD FONT.

Model	Radial	Spiral	# Parameters
PDNet no DCp	27.02 / 0.6747	28.02 / 0.6946	156k
Adjoint + DCp	27.11 / 0.6471	31.70 / 0.7213	0
DIP	29.57 / 0.5148	29.79 / 0.5249	0
PDNet w gridding	31.12 / 0.6887	31.45 / 0.7126	156k
U-net on Adjoint + DCp	32.26 / 0.7224	32.82 / 0.7460	481k
NC-PDNet	32.66 / 0.7327	33.08 / 0.7534	156k

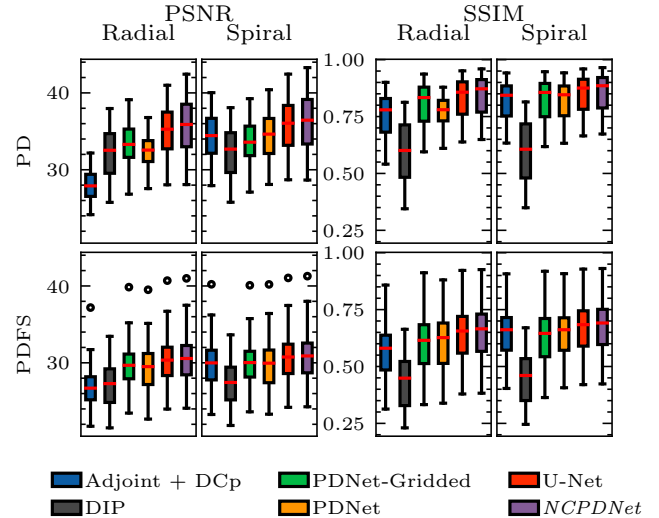


Fig. 3. **Single-coil knee dataset**: Quantitative results of the different networks in the single-coil setting for both fastMRI contrasts.

chose to not run the DIP experiments on brain data in order to keep the amount of computation reasonable.

Third, we tested the trained networks on an $AF = 8$, i.e. two times larger than the acceleration factor they were confronted in training.

d) *Computational efficiency*: The time spent for image reconstruction by a given method is also an important factor to take into account when choosing which algorithm to be used in the clinical realm. Indeed, in clinical applications we need to visualize the reconstructed MR images prior to the end of the exam to permit any potential re-start of the pulse MR sequence in case of the presence of significant artifacts impeding accurate image-based diagnosis. Additionally, this allows the physician to report on the scans directly to the patient once the exam is completed. For that reason, we decided to report the reconstruction times associated with the different methods and networks.

B. Quantitative results

a) *2D single-coil fastMRI dataset*: The quantitative results in Tab. I show that *NC-PDNet* outperforms the concurrent approaches. Moreover, these results outline that combining unrolling and DCp with an accurate NDFT approximation (i.e. the NUFFT) within the same deep learning architecture is instrumental in non-Cartesian MR image reconstruction as using only one of these three ingredients will lead to degraded performance. It is worth mentioning that DIP performs

TABLE II

MEAN PSNR / SSIM ON THE VALIDATION VOLUMES OF THE DIFFERENT APPROACHES AVERAGED OVER BOTH CONTRASTS (KNEE FASTMRI) IN THE MULTI-COIL SETTING. THE BEST RESULTS ARE IN BOLD FONT.

Model	Radial	Spiral	# Parameters
Adjoint + DCp	25.91 / 0.6486	31.36 / 0.7197	0
DIP	29.21 / 0.5834	29.19 / 0.5832	0
U-net on Adjoint + DCp	38.78 / 0.9106	40.02 / 0.9215	481k
NC-PDNet	40.00 / 0.9191	40.68 / 0.9255	163k

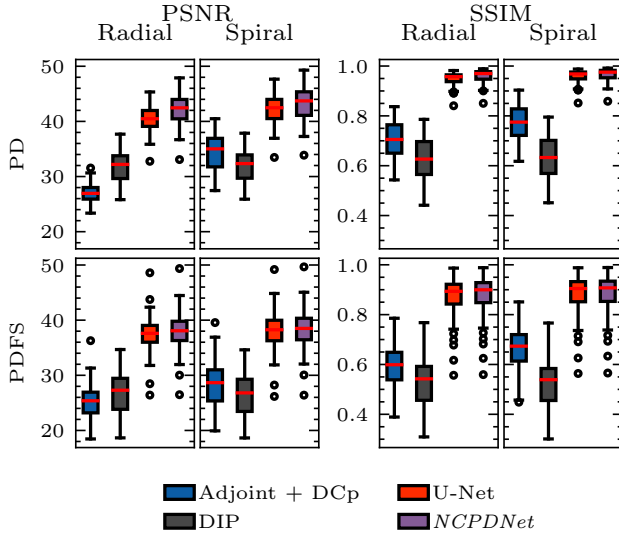


Fig. 4. Multi-coil knee dataset: Quantitative results of the different networks in the multi-coil setting for both fastMRI contrasts.

poorly compared to U-net. Additionally, the quantitative scores are slightly improved for spiral readout compared to radial whatever the reconstruction strategy, likely because this under-sampling scheme has a better coverage of high frequencies in the k-space.

b) *2D multi-coil fastMRI dataset*: In regards to the 2D multi-coil numerical experiments, we selected only the two best performers of the single-coil studies for comparison purposes, along with the adjoint baseline for reference and DIP being a strong contender for the generalization studies. As reported in Tab. II the Adjoint with DCp has a very low performance for the radial trajectory, hence the U-net working on top of this raw strategy is not able to correct sufficiently well for aliasing artifacts in the image domain in order to be competitive with the *NC-PDNet*. This illustrates that *NC-PDNet* has the ability to overcome situations where the trajectory might be causing issues to naive methods. We also observed for the *NC-PDNet* that the PSNR improvement in spiral imaging compared to radial imaging is larger in the multi-coil setting. Finally, it can be noted that the sensitivity maps refinement has a relatively low impact on the number of parameters in the *NC-PDNet*, with only 7k additional parameters in θ_r (163k vs 156k parameters in bottom rows of Tab. II vs Tab. I).

c) *3D single-coil brain OASIS dataset*: We can see in Tab. III that the *NC-PDNet* outperforms the baseline models on the OASIS dataset by a significant margin, even though its size was reduced due to memory constraints. We specify that the SSIM was computed on slices along the first dimension of the

TABLE III

MEAN PSNR / SSIM ON THE VALIDATION VOLUMES OF THE DIFFERENT APPROACHES FOR THE OASIS DATASET (3D SETTING). THE BEST RESULTS ARE IN BOLD FONT.

Model	Radial	#Parameters
Adjoint + DCp	27.51 / 0.5183	0
U-net on Adjoint + DCp	31.42 / 0.8432	1.6M
NC-PDNet	33.76 / 0.9160	67k




Fig. 5. 3D OASIS brain dataset: PSNR and SSIM distribution of the different networks in the 3D radial under-sampling scenario.

volumes in the 3D case. This choice was arbitrary.

TABLE IV

MEAN PSNR/SSIM ON THE VALIDATION VOLUMES OF THE DIFFERENT APPROACHES AVERAGED OVER BOTH CONTRASTS (KNEE FASTMRI) IN THE MULTI-COIL REVERSE SETTING. BEST RESULTS ARE IN BOLD FONT.

Model	Radial	Spiral	# Parameters
Adjoint + DCp	25.91 / 0.6486	31.36 / 0.7197	0
DIP	29.21 / 0.5834	29.19 / 0.5832	0
U-net on Adjoint + DCp (trained on different trajectory)	27.94 / 0.87	26.35 / 0.8850	481k
NC-PDNet (trained on different trajectory)	37.86 / 0.9079	36.28 / 0.9052	163k

d) *Reverse trajectories in 2D multi-coil knee imaging*: The quantitative results for the reverse setting shown in Tab. IV reveal contrasting results. The drop in PSNR is significant for both networks (compare Tab. IV line by line with the corresponding reference in Tab. II): The weaker performance is achieved for the U-net trained on radial and evaluated on spiral (by more than 13dB), while the drop for *NC-PDNet* is limited to approximately 4.5dB in the same use case and even less in the radial validation case (37.86dB vs 40dB). In contrast, the situation in SSIM is very much controlled in the *NC-PDNet*'s case (a bit less so for the U-net).

e) *2D multi-coil brain fastMRI dataset*: The quantitative results for the brain dataset are not to be compared head-to-head to the results obtained for the knee data. Indeed there is a significant mismatch in reconstruction difficulty, the brain data being easier to reconstruct than the knee data metrics-wise, based on the fastMRI public leaderboard⁷. However, in Tab. V we can see that the neural networks still largely outperform the baseline (Adjoint + DCp) for both trajectories. We also observed that the *NC-PDNet* outperforms the U-net

⁷<https://fastmri.org/leaderboards/>

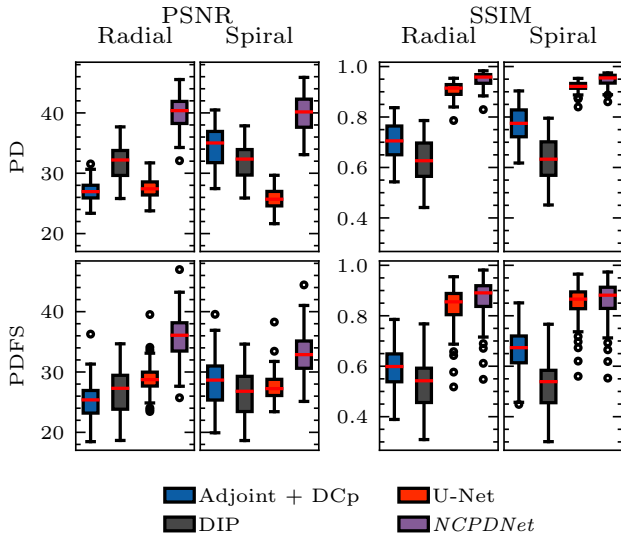


Fig. 6. **Multi-coil knee dataset (reverse setting)**: Quantitative results of the different networks in the multi-coil reverse setting for both fastMRI contrasts.

on both trajectories, this time with a bigger gap, suggesting that it generalizes in a better way.

TABLE V

MEAN PSNR/SSIM ON THE VALIDATION VOLUMES OF THE DIFFERENT APPROACHES AVERAGED OVER ALL BRAIN FASTMRI IMAGING CONTRASTS IN THE MULTI-COIL SETTING.

Model	Radial	Spiral	# Parameters
Adjoint + DCp	27.31 / 0.6028	32.23 / 0.6603	0
U-net on Adjoint + DCp (trained on knee data)	37.88 / 0.9234	38.76 / 0.9302	481k
NC-PDNet (trained on knee data)	39.48 / 0.9368	39.81 / 0.9390	163k

f) *Higher acceleration factor*: The quantitative results in Tab. VI for a higher AF in validation, namely $AF = 8$ compared to $AF = 4$ in training, allow us to draw the same conclusions as the previous generalization experiments: the design of *NC-PDNet* is instrumental in getting more robust results for different trajectories.

TABLE VI

MEAN PSNR / SSIM ON THE VALIDATION VOLUMES OF THE DIFFERENT APPROACHES FOR BOTH CONTRASTS (KNEE FASTMRI) IN THE MULTI-COIL SETTING FOR $AF = 8$.

Model	Radial	Spiral
Adjoint + DCp	25.62 / 0.6097	30.00 / 0.6686
DIP	29.08 / 0.5782	29.21 / 0.5841
U-net on Adjoint + DCp	34.11 / 0.8592	34.60 / 0.8665
NC-PDNet	36.71 / 0.8887	37.47 / 0.8967

g) *Reconstruction times*: The reconstruction times given in Tab. VII show that the use of the NUFFT in the unrolled setting induces a 4.5-fold to 9-fold increase in execution time for 2D and 3D image reconstruction⁸. However, the use of the DCp merely increases the reconstruction time by only 33ms.

⁸The reconstruction times observed for 3D are much less than those in 2D because of the different image sizes as well as smaller sizes of the networks.

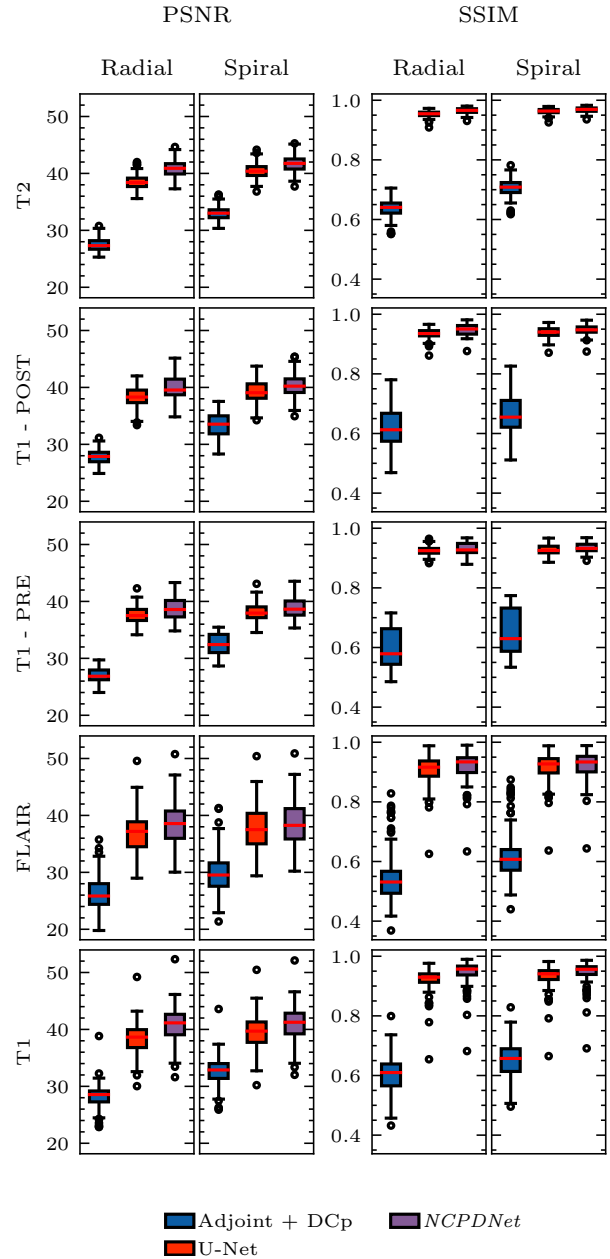


Fig. 7. **2D multi-coil fastMRI brain dataset**: Quantitative results of the different networks in the brain multi-coil setting for the multiple contrasts (T1, T2, FLAIR) with some distinction on T1-w imaging related to Gadolinium injection (T1-PRE, T1-POST).

A number that stands out is of course the reconstruction time for DIP. This was already noted in [62], where these authors proposed an acceleration strategy that we implemented. However the computation time remains unrealistic for clinical applications.

C. Qualitative results

For all qualitative results, the images were selected at random and do not necessarily have the same contrast. For each setting, the missing acquisition trajectory's reconstruction figures (if any) can be found in the Supplementary Material in Figs. S2-S4.

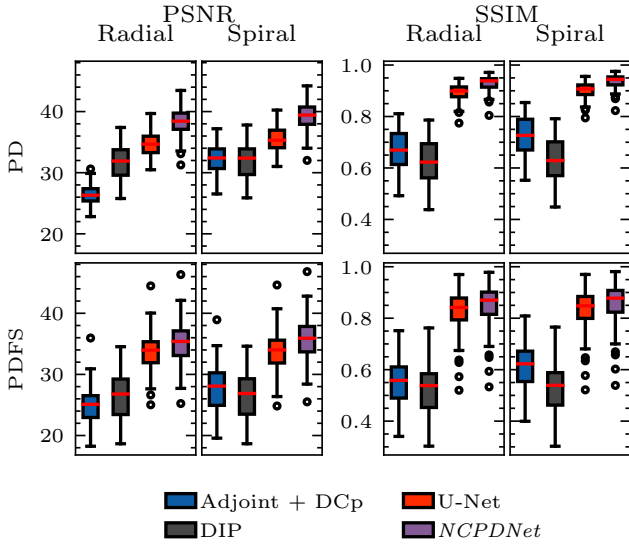


Fig. 8. **2D multi-coil** ($AF = 8$): Quantitative results of the different networks in the multi-coil setting for both fastMRI contrasts (knee dataset) with a higher acceleration factor during validation ($AF = 8$) compared to training ($AF = 4$).

TABLE VII

RECONSTRUCTION TIMES OF A SINGLE SLICE IN MILLISECONDS FOR THE DIFFERENT NETWORKS IN THE DIFFERENT ACQUISITION SCENARIOS BASED ON 2D/3D RADIAL UNDER-SAMPLING.

Model	Single-coil 2D	Multi-coil 2D	3D
PDNet no DCp	446	NA	NA
Adjoint + DCp	101	135	7
DIP	110k	133k	NA
PDNet w gridding	112	NA	NA
U-net on Adjoint + DCp	110	145	9
NC-PDNet	479	661	80

We chose to highlight some key parts of the reconstruction to pay attention to for the reader to better understand the difference between the U-net reconstruction and ours.

a) *2D single-coil knee fastMRI dataset*: The visual inspection of the reconstructed MR images confirm the quantitative measurements. In particular, one can visualize in Fig. 9 that the image’s inner contrast is better recovered by the *NC-PDNet* and U-net, and the structures are sharper for both architectures compared to the other competitors. The zooms allow us to identify the blurriness in DIP, the benefit of using DCp in the *NC-PDNet* architecture compared to the gridded and vanilla versions of PDNet.

b) *2D multi-coil knee fastMRI dataset*: The qualitative results for the multi-coil setting, available in Fig. 10 confirm the quantitative results. We see on the right part of the knee (cf. red frame and zooms) that the reconstruction by the U-net after the application of the adjoint with DCp is not completely faithful. The reconstruction by the *NC-PDNet* is much more representative of the ground truth anatomy. In regards to the DIP solution, it appears very blurry compared to U-Net and *NC-PDNet* and hence this confirms the results we obtained in the single-coil setting.

c) *3D single-coil brain imaging (OASIS dataset)*: Similarly to previous numerical experiments, the qualitative results shown

in Fig. 11 confirm the quantitative scores for the 3D brain imaging setting. Although the *NC-PDNet*’s reconstruction is blurred compared to the ground truth, we notice a much better resolution than the baselines as can be seen in the magnified views.

d) *Reverse trajectories in 2D multi-coil knee imaging*: The qualitative results for the reverse setting, shown in Figs. 12-13, allow us to look at the quantitative metrics from a different perspective. We actually observed that the reconstructions are excellent in terms of image quality and that the degradation compared to the regular setting is minimally visible (except the radial artifacts for the *NC-PDNet* trained on spiral acquisition and evaluated on radial spokes in the right part of the knee). This confirms that the SSIM is the best metric to monitor in order for measuring success in generalization to other under-sampling patterns.

e) *2D multi-coil brain fastMRI dataset*: The qualitative results shown in Fig. 14 for the 2D multi-coil brain FLAIR image confirm that the *NC-PDNet* was able to better generalize than the U-net.

f) *Higher acceleration factor*: We see in Fig. 15 that the higher AF has a very severe effect on all the reconstructions in terms of blur. Although the image based on the *NC-PDNet* architecture also suffers from a loss in resolution, it remains the best reconstruction with the least amount of blur and the best recovery of finer details.

VI. DISCUSSION AND CONCLUSION

In this work, we demonstrated how to apply the framework of unrolled neural networks to the problem of non-Cartesian MRI reconstruction. In particular, we showed via an ablation study, the importance of using the mechanism of density compensation [24] for this setting and how it is instrumental in obtaining the best possible results in a deep learning framework. We also managed to show how this mechanism can be extended to challenging acquisition scenarios like multi-coil and 3D reconstruction with very minor adjustments. We conducted additional generalizability experiments to measure the robustness of the trained networks to out-of-distribution settings, and concluded that the visual degradation is minimal for all networks, and that additionally the *NC-PDNet* is quantitatively consistent. In particular, we found out that the generalization to other organs is possible even when using a single organ in training. Finally the modular design of the code allows implementation and improvement of the results.

Some issues still remain open regarding this work. The first is, of course, that of transferability/generalizability. Indeed, these networks have been trained on a specific dataset with given experimental conditions, a given MR system vendor (Siemens-Healthineers, Erlangen, Germany), and a specific acquisition setup and organ. The question of generalizability from one vendor to another has been partially answered in [14], where it has been observed that the results degrade when the networks trained with one vendor are applied to another. Although we partially addressed the concerns regarding the acquisition setup and the organs with our out-of-distribution experiments, a more systematic study is still required.




Fig. 9. 2D single-coil radial acquisition (knee fastMRI dataset, PD contrast): Reconstruction results for a specific slice (16th slice of file1001184, part of the validation set). The top row represents the reconstruction using the different methods, while the bottom one represents the zoom highlighted by a red frame in the top row. The volume-wise PSNR and SSIM scores are shown on the top of each full FOV image. The arrow indicates a streaking artefact.




Fig. 10. 2D multi-coil radial acquisition (knee fastMRI dataset, PDFS contrast): Reconstruction results for a specific slice (16th slice of file1000000, part of the validation set). The top row represents the reconstruction using the different methods, while the bottom one represents the zoom highlighted by a red frame in the top row. The volume-wise PSNR and SSIM scores are shown on the top of each full FOV image. The arrow indicates a strong streaking artefact.

A more challenging question that can be raised is whether a lack of generalizability is intrinsic to the training process or not: The need for training data generally translates to using retrospective under-sampling to be able to learn how to correct for aliasing artifacts from a “correct” ground truth. How do the networks generalize to prospective under-sampling, especially in non-Cartesian acquisition scenarios where off-resonance effects come into play?

This question is very difficult to answer systematically. In the case of prospective acquisition, as no ground truth is available, the methods cannot be compared one another. The only viable option is to have qualitative results, such as those available in [63], which showed in a specific case that an unrolled neural network was robust enough to reconstruct prospectively under-sampled Cartesian data.

One could try to simulate the non-idealities of the clinical

setting (such as B0 field inhomogeneities, gradient inaccuracies/delays, eddy currents, etc...) as was done in [64], but the gain and correct implementation of such simulations remains an open question. A final issue on transferability arises in this work for the 3D data, which consisted of magnitude volumes (and not of raw k-space data as in the 2D case). How can the networks generalize to complex-valued data? There exists ways to generate an artificial phase information, and they could be of help in order to be closer to the real use-case [64]. Relevant data augmentation mechanisms [65] can be one of these solutions.

One disadvantage of the current implementation of unrolled networks in the non-Cartesian setting is their slowness compared to other baseline networks. This has been diagnosed in this work to be due to the NUFFT. Recent work has been carried out to implement an efficient NUFFT on the




Fig. 11. 3D radial acquisition (OASIS dataset, T1 contrast): Reconstruction results for a specific slice (101st slice of sub-OAS30001_ses-d0129_run-01_T1w, part of the validation set). The top row represents the reconstruction using the different methods, while the bottom one represents the zoom highlighted by a red frame in the top row. The volume-wise PSNR and SSIM scores are shown on the top of each full FOV image. The arrow indicates a blurred out structure.




Fig. 12. 2D multi-coil radial acquisition (reverse setting, knee dataset): Reconstruction results for a specific slice (16th slice of file1000000, part of the validation set) with networks trained with spiral trajectories. The top row represents the reconstruction using the different methods, while the bottom one represents the zoom highlighted by a red square in the top row. The volume-wise PSNR and SSIM scores are shown on the top of each full FOV image. The arrow indicates a strong streaking artefact.

GPU, cuFINUFFT [66], although not directly in a framework allowing auto-differentiation, making its use at training time difficult. However, fast implementations could very well be embedded in the network at reconstruction allowing us to alleviate the slowness of NUFFT in auto-differentiation frameworks. Complementary to that, a recent work [67] also tackles the problem of the slowness of the NUFFT by introducing a TPU implementation.

Finally, the 3D setting illustrated a major blocker on the road ahead for unrolled neural networks: Memory requirements for the backpropagation during training. Indeed, while the *NC-PDNet* outperformed the baseline models, none of the networks seemed to overfit, implying that there is room for networks with larger capacity. However, training such networks in the current state with a basic implementation on regular 32GB GPUs is out-of-reach. Some possible workarounds of

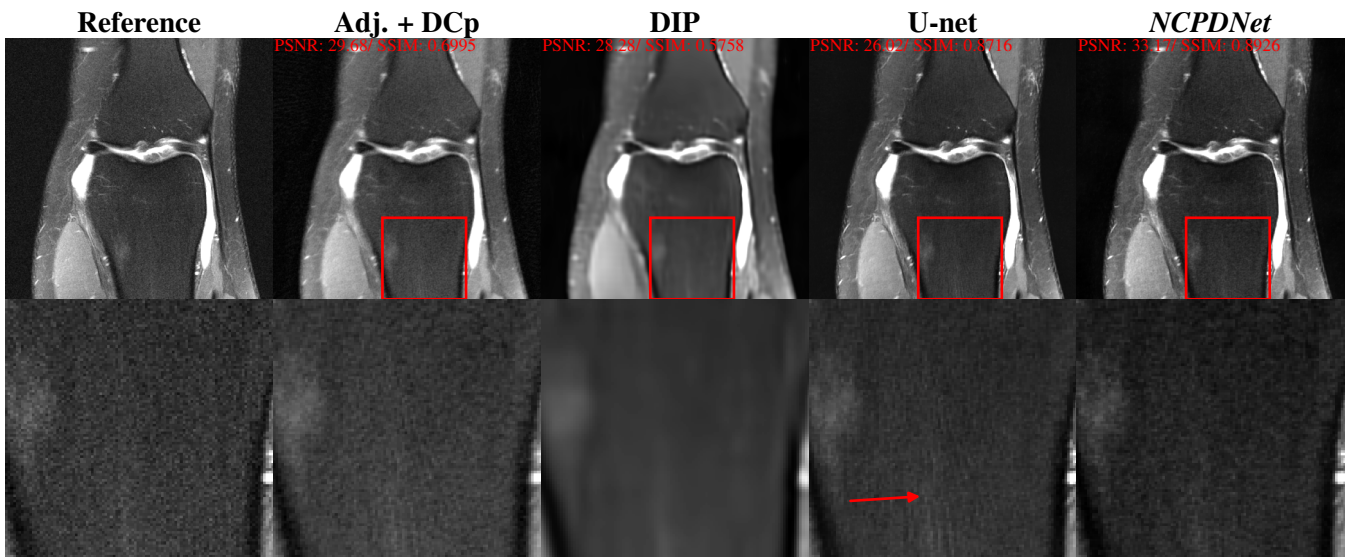


Fig. 13. 2D multi-coil spiral acquisition (reverse setting, knee dataset): Reconstruction results for a specific slice (16th slice of `file1000000`, part of the validation set) with networks trained on radial trajectories. The top row represents the reconstruction using the different methods, while the bottom one represents the zoom highlighted by a red frame in the top row. The volume-wise PSNR and SSIM scores are shown on the top of each full FOV image. The arrow indicates a light streaking artefact.




Fig. 14. 2D multi-coil radial acquisition (brain fastMRI dataset, FLAIR contrast): Reconstruction results for a specific slice (6th slice of `file_brain_AXFLAIR_200_6002447`) from the brain fastMRI dataset with networks trained on knee data. The top row represents the reconstruction using the different methods, while the bottom one represents the zoom highlighted by a red frame in the top row. The volume-wise PSNR and SSIM scores are shown on the top of each full FOV image. The arrow indicates a non-recovered structure.

this problem exist on the implementation side, using multiple GPUs, like model parallelism [68] or pipelining [69]. Some other solutions exist on the modeling side, by using networks that either reduce the memory requirements of the backpropagation like gradient checkpointing [70], invertible networks [49], [71] or by getting rid of the backpropagation altogether like implicit deep learning [72]–[74].

ACKNOWLEDGMENTS

We acknowledge the financial support of the Cross-Disciplinary Program on Numerical Simulation of CEA for the project entitled SILICOSMIC. This work was granted access to the HPC resources of IDRIS under the allocation 2021-AD011011153 made by GENCI.




Fig. 15. 2D multi-coil radial acquisition (knee fastMRI dataset, PDFS contrast) $AF = 8$: Reconstruction results for a specific slice (16th slice of file1000000, part of the validation set) for an AF 8. The top row represents the reconstruction using the different methods, while the bottom one represents the zoom highlighted by a red frame in the top row. The arrow indicates a strong streaking artefact.

REFERENCES

- [1] Z. Ramzi, P. Ciuciu, and J.-L. Starck, "Density compensated unrolled networks for non-Cartesian MRI reconstruction," in *18th IEEE International Symposium on Biomedical Imaging (ISBI)*, 2021.
- [2] K. P. Pruessmann *et al.*, "SENSE: Sensitivity encoding for fast MRI," *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 11 1999.
- [3] M. A. Griswold *et al.*, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 6 2002.
- [4] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [5] M. Guerquin-Kern *et al.*, "A Fast Wavelet-Based Reconstruction Method for Magnetic Resonance Imaging," *IEEE Transactions on Medical Imaging*, vol. 30, no. 9, pp. 1649–1660, 2011.
- [6] L. Chaïri *et al.*, "A wavelet-based regularized reconstruction algorithm for sense parallel MRI with applications to neuroimaging," *Medical image analysis*, vol. 15, no. 2, pp. 185–201, 2011.
- [7] J. A. Fessler and B. P. Sutton, "Nonuniform Fast Fourier Transforms Using Min-Max Interpolation," *IEEE Transactions on Signal Processing*, vol. 51, no. 2, pp. 560–74, 2003.
- [8] J. Caballero *et al.*, "Dictionary learning and time sparsity for dynamic MR data reconstruction," *IEEE Transactions on Medical Imaging*, vol. 33, no. 4, pp. 979–994, 2014.
- [9] K. Gregor and Y. Lecun, "Learning fast approximations of sparse coding," in *Proc. of the 27th International Conference on Machine Learning (ICML)*, 2010.
- [10] B. Zhu *et al.*, "Image reconstruction by domain-transform manifold learning," *Nature*, vol. 555, no. 7697, pp. 487–492, 3 2018.
- [11] T. Eo *et al.*, "KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images," *Magnetic Resonance in Medicine*, vol. 80, no. 5, pp. 2188–2201, 2018.
- [12] K. Hammernik *et al.*, "Learning a Variational Network for Reconstruction of Accelerated MRI Data," *Magnetic Resonance in Medicine*, vol. 3071, no. 79, pp. 3055–3071, 2018.
- [13] J. Schlemper *et al.*, "A deep cascade of convolutional neural networks for MR image reconstruction," *IEEE Transactions on Medical Imaging*, vol. 37, no. 2, pp. 491–503, 2018.
- [14] M. J. Muckley *et al.*, "Results of the 2020 fastMRI Challenge for Machine Learning MR Image Reconstruction," *IEEE Transactions on Medical Imaging*, May 2021.
- [15] D. C. Peters *et al.*, "Undersampled projection reconstruction applied to MR angiography," *Magnetic Resonance in Medicine*, vol. 43, no. 1, pp. 91–101, 2000.
- [16] M. A. Bernstein, K. F. King, and X. J. Zhou, "Advanced pulse sequence techniques," in *Handbook of MRI Pulse Sequences*. Elsevier, 2004, pp. 802–954.
- [17] L. Feng *et al.*, "Golden-angle radial sparse parallel MRI: combination of compressed sensing, parallel imaging, and golden-angle radial sampling for fast and flexible dynamic volumetric MRI," *Magnetic resonance in Medicine*, vol. 72, no. 3, pp. 707–717, 2014.
- [18] D. G. Nishimura, P. Irarrazabal, and C. H. Meyer, "A velocity k-space analysis of flow effects in echo-planar and spiral imaging," *Magnetic resonance in Medicine*, vol. 33, no. 4, pp. 549–556, 1995.
- [19] K. T. Block *et al.*, "Towards Routine Clinical Use of Radial Stack-of-Stars 3D Gradient-Echo Sequences for Reducing Motion Sensitivity," *Journal of the Korean Society of Magnetic Resonance in Medicine*, vol. 18, no. 2, p. 87, 2014.
- [20] D. R. Thedens *et al.*, "Fast magnetic resonance coronary angiography with a three-dimensional stack of spirals trajectory," *Magnetic Resonance in Medicine*, vol. 41, no. 6, pp. 1170–1179, 1999.
- [21] C. Lazarus *et al.*, "3D variable-density SPARKLING trajectories for high-resolution T2*-weighted magnetic resonance imaging," *NMR in Biomedicine*, vol. 33, no. 9, p. e4349, 2020.
- [22] C. G. R. *et al.*, "Optimizing full 3D SPARKLING trajectories for high-resolution T2*-weighted magnetic resonance imaging," CEA/NeuroSpin, Université Paris-Saclay, Inria Parietal, Gif-sur-Yvette, France, in revision to *IEEE Transactions on Medical Imaging*, Aug. 2021. [Online]. Available: <https://hal.inria.fr/hal-03090471>
- [23] M. Abadi *et al.*, "TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems," Tech. Rep., 2016. [Online]. Available: <http://arxiv.org/abs/1603.04467>
- [24] J. G. Pipe and P. Menon, "Sampling density compensation in MRI: Rationale and an iterative numerical solution," *Magnetic Resonance in Medicine*, vol. 41, no. 1, pp. 179–186, 1999.
- [25] J. D. Trzasko *et al.*, "A Preconditioned ADMM Strategy for Field-Corrected Non-Cartesian MRI Reconstruction," p. 1535, 2014.
- [26] F. Ong, M. Uecker, and M. Lustig, "Accelerating Non-Cartesian MRI Reconstruction Convergence Using k-Space Preconditioning," *IEEE Transactions on Medical Imaging*, vol. 39, no. 5, pp. 1646–1654, 2020.
- [27] J. Adler and O. Öktem, "Learned primal-dual reconstruction," *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1322–1332, 2018.
- [28] 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, 2019.
- [29] J. Schlemper *et al.*, "Nonuniform variational network: Deep learning for accelerated nonuniform MR image reconstruction," pp. 57–64, 2019.
- [30] K. L. Wright *et al.*, "Non-Cartesian parallel imaging reconstruction," *Journal of Magnetic Resonance Imaging*, vol. 40, no. 5, pp. 1022–1040, 2014.
- [31] S. Poddar and M. Jacob, "Dynamic MRI Using Smoothness Regularization on Manifolds (SToRM)," *IEEE Transactions on Medical Imaging*, vol. 35, no. 4, pp. 1106–1115, 2016.
- [32] S. Poddar *et al.*, "Manifold Recovery Using Kernel Low-Rank Regularization: Application to Dynamic Imaging," *IEEE Transactions on Computational Imaging*, vol. 5, no. 3, pp. 478–491, 2019.
- [33] U. Nakarmi *et al.*, "M-MRI: A manifold-based framework to highly accelerated dynamic magnetic resonance imaging," *14th IEEE International Symposium on Biomedical Imaging (ISBI)*, pp. 19–22, 2017.
- [34] A. H. Ahmed *et al.*, "Free-breathing and ungated dynamic mri using navigator-less spiral storm," *IEEE Transactions on Medical Imaging*, vol. 39, no. 12, pp. 3933–3943, 2020.
- [35] F. Knoll *et al.*, "Advancing machine learning for MR image reconstruction with an open competition: Overview of the 2019 fastMRI challenge," *Magnetic Resonance in Medicine*, vol. 84, no. 6, pp. 3054–3070, 2020.
- [36] Y. Yang *et al.*, "Deep ADMM-net for compressive sensing MRI," in *Neural Information Processing Systems*, 2016, pp. 10–18.
- [37] D. Gilton, G. Ongie, and R. Willett, "Neumann Networks for Linear Inverse Problems in Imaging," *IEEE Transactions on Computational Imaging*, vol. 6, pp. 328–343, 2019.
- [38] Z. Ramzi, J.-L. Starck, and P. Ciuciu, "XPDNet for MRI Reconstruction: an application to the 2020 fastMRI challenge," May 2021.
- [39] N. Chauffert *et al.*, "Variable density sampling with continuous trajectories. Application to MRI," *SIAM Imaging Sciences*, vol. 7, no. 4, pp. 1962–1992, Nov. 2014.
- [40] B. Adcock *et al.*, "Breaking the coherence barrier: A new theory for compressed sensing," in *Forum of Math., Sigma*, vol. 5. Cambridge University Press, 2017.
- [41] C. Boyer *et al.*, "On the generation of sampling schemes for MRI," *SIAM Imaging Sciences*, vol. 9, no. 4, pp. 2039–2072, 2016.
- [42] D. Shen *et al.*, "Rapid reconstruction of highly undersampled, non-Cartesian real-time cine k-space data using a perceptual complex neural network (PCNN)," *NMR in Biomedicine*, vol. 34, no. 1, pp. 1–12, 2021.
- [43] M. J. Muckley *et al.*, "TorchKbNufft: A high-level, hardware-agnostic non-uniform fast fourier transform," in *ISMRM Workshop on Data Sampling & Image Reconstruction*, 2020.
- [44] M. O. Malavé *et al.*, "Reconstruction of undersampled 3D non-Cartesian image-based navigators for coronary MRA using an unrolled deep learning model," *Magnetic Resonance in Medicine*, vol. 84, no. 2, pp. 800–812, 8 2020.
- [45] P. Putzky and M. Welling, "Invert to Learn to Invert," in *Advances in neural information processing systems*, 2019.
- [46] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 9446–9454.
- [47] J. Yoo *et al.*, "Time-Dependent Deep Image Prior for Dynamic MRI," *IEEE Transactions on Medical Imaging*, 2021.
- [48] T. Küstner *et al.*, "CINENet: deep learning-based 3D cardiac CINE MRI reconstruction with multi-coil complex-valued 4D spatio-temporal convolutions," *Scientific Reports*, vol. 10, no. 1, pp. 1–13, 2020.
- [49] M. Kellman *et al.*, "Memory-efficient learning for large-scale computational imaging," *IEEE Transactions on Computational Imaging*, vol. 6, pp. 1403–1414, 2020.
- [50] T. Meinhardt *et al.*, "Learning proximal operators: Using denoising networks for regularizing inverse imaging problems," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 1781–1790.
- [51] J. Adler and O. Öktem, "Solving ill-posed inverse problems using iterative deep neural networks," *Inverse Problems*, vol. 33, no. 12, pp. 1–24, 2017.
- [52] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," pp. 234–241, 2015.

- [53] A. Sriram *et al.*, “End-to-end variational networks for accelerated MRI reconstruction,” in *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2020.
- [54] M. Uecker *et al.*, “ESPIRiT-An eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA,” *Magnetic Resonance in Medicine*, vol. 71, no. 3, pp. 990–1001, 2014.
- [55] E. K. Cole, J. Pauly, and J. Cheng, “Complex-valued convolutional neural networks for MRI reconstruction,” Montreal, Canada, p. 4714, 2019.
- [56] P. J. Beatty, D. G. Nishimura, and J. M. Pauly, “Rapid gridding reconstruction with a minimal oversampling ratio,” *IEEE Transactions on Medical Imaging*, vol. 24, no. 6, pp. 799–808, 2005.
- [57] J. Zbontar *et al.*, “fastMRI: An Open Dataset and Benchmarks for Accelerated MRI,” Tech. Rep., 2018. [Online]. Available: <https://arxiv.org/pdf/1811.08839.pdf>
- [58] P. J. LaMontagne *et al.*, “OASIS-3: Longitudinal Neuroimaging, Clinical, and Cognitive Dataset for Normal Aging and Alzheimer Disease,” Tech. Rep., 2019. [Online]. Available: <https://www.medrxiv.org/content/10.1101/2019.12.13.19014902v1>
- [59] C. M. Sandino *et al.*, “Compressed sensing: From research to clinical practice with deep neural networks,” *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 117–127, 2020.
- [60] K. Epperson *et al.*, “Creation of fully sampled MR data repository for compressed sensing of the knee purpose,” 2013. [Online]. Available: <http://mridata.org/list?project=StanfordFullysampled3DFSEKnees>
- [61] Z. Wang *et al.*, “Image quality assessment: From error visibility to structural similarity,” *IEEE Transactions on Image Processing*, vol. 13, no. 4, 2004.
- [62] M. Z. Darestani and R. Heckel, “Accelerated MRI with Un-trained Neural Networks,” 2020. [Online]. Available: <http://arxiv.org/abs/2007.02471>
- [63] Z. Ramzi *et al.*, “Is good old GRAPPA dead?” May 2021.
- [64] J. Schlemper, M. Salehi, and M. Sofka, “Mutli-coil Magnetic Resonance Imaging Using Deep Learning,” 2020. [Online]. Available: <https://patents.google.com/patent/US20200294287A1/en>
- [65] Z. Fabian, R. Heckel, and M. Soltanolkotabi, “Data augmentation for deep learning based accelerated MRI reconstruction with limited data,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 3057–3067.
- [66] Y.-h. Shih, J. Blaschke, and A. H. Barnett, “cuFINUFFT : a load-balanced GPU library for general-purpose nonuniform FFTs,” Tech. Rep., 2021. [Online]. Available: <https://arxiv.org/abs/2102.08463>
- [67] T. Lu, T. Marin *et al.*, “Nonuniform fast Fourier transform on TPUs,” in *18th IEEE International Symposium on Biomedical Imaging (ISBI)*, 2021, pp. 783–787.
- [68] N. Shazeer *et al.*, “Mesh-tensorflow: Deep learning for supercomputers,” in *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [69] Y. Huang *et al.*, “GPipe: Efficient training of giant neural networks using pipeline parallelism,” in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [70] T. Chen *et al.*, “Training deep nets with sublinear memory cost,” Tech. Rep., 2016. [Online]. Available: <http://arxiv.org/abs/1604.06174>
- [71] M. E. Sander *et al.*, “Momentum Residual Neural Networks,” in *ICML*, 2021. [Online]. Available: <http://arxiv.org/abs/2102.07870>
- [72] S. Bai, J. Zico Kolter, and V. Koltun, “Deep equilibrium models,” in *Advances in Neural Information Processing Systems*, 2019.
- [73] D. Gilton, G. Ongie, and R. Willett, “Deep equilibrium architectures for inverse problems in imaging,” Tech. Rep., 2021. [Online]. Available: <http://arxiv.org/abs/2102.07944>
- [74] Z. Ramzi *et al.*, “SHINE: SHaring the INverse Estimate from the forward pass for bi-level optimization and implicit models,” arXiv preprint arXiv:2106.00553, May 2021.