

# GrappaNet: Combining Parallel Imaging with Deep Learning for Multi-Coil MRI Reconstruction

Anuroop Sriram<sup>1</sup>    Jure Zbontar<sup>1</sup>    Tullie Murrell<sup>1</sup>    C. Lawrence Zitnick<sup>1</sup>  
 Aaron Defazio<sup>1</sup>    Daniel K. Sodickson<sup>2</sup>

<sup>1</sup>Facebook AI Research (FAIR)    <sup>2</sup>NYU School of Medicine

## Abstract

*Magnetic Resonance Image (MRI) acquisition is an inherently slow process which has spurred the development of two different acceleration methods: acquiring multiple correlated samples simultaneously (parallel imaging) and acquiring fewer samples than necessary for traditional signal processing methods (compressed sensing). Both methods provide complementary approaches to accelerating MRI acquisition.*

*In this paper, we present a novel method to integrate traditional parallel imaging methods into deep neural networks that is able to generate high quality reconstructions even for high acceleration factors. The proposed method, called GrappaNet, performs progressive reconstruction by first mapping the reconstruction problem to a simpler one that can be solved by a traditional parallel imaging method using a neural network, followed by an application of a parallel imaging method, and finally fine-tuning the output with another neural network. The entire network can be trained end-to-end. We present experimental results on the recently released fastMRI dataset [20] and show that GrappaNet can generate higher quality reconstructions than competing methods for both 4× and 8× acceleration.*

## 1. Introduction

Magnetic Resonance Imaging (MRI) is the leading diagnostic modality for a wide range of disorders including musculoskeletal, neurological, and oncological diseases. However, the physics of the MRI data acquisition process make it inherently slower than alternate modalities like CT or X-Ray. As a consequence, increasing the speed of MRI acquisition has been a major ongoing research goal for decades.

Parallel Imaging (PI) is one of the most important and successful developments in reducing MRI scan time [3, 12]. The technique requires the use of multiple physical receiver coils to simultaneously record different views of the object being imaged. Parallel imaging is the default option for

many scan protocols and it is supported by almost all modern clinical MRI scanners.

Another approach to accelerating MR imaging is the use of Compressed Sensing (CS), which can speed up MRI acquisition by acquiring fewer samples than required by traditional signal processing methods. To overcome aliasing artifacts introduced by violating the Shannon-Nyquist sampling theorem, CS methods incorporate additional a priori knowledge about the images. Recently, the use of learned image priors through the use of deep learning have rapidly gained in popularity [4, 15, 18, 21]. These approaches have shown a significant improvement in image reconstruction quality, particularly for non-parallel MRIs.

In this paper, we show that a novel combination of classical parallel imaging techniques with deep neural networks can achieve higher acceleration factors than using either approach alone. Utilizing parallel imaging in deep learning approaches to reconstruction is challenging. The relation between the captured views changes for each scan and is dependent on the configuration of the detectors with respect to the object being imaged.

To address this challenge we introduce GrappaNet, a new neural network architecture that incorporates parallel imaging. GrappaNet contains a GRAPPA layer that learns a scan-specific reconstruction function to combine the views captured during parallel imaging. To allow the network to fully utilize all the information captured during parallel imaging, the reconstruction is performed jointly across all the complex-valued views captured during the parallel imaging process. Unlike many previous approaches [4], the views are not combined until the final layer to produce the output reconstruction. The model uses a progressive refinement approach in both k-space (frequency domain) and image space to both aid in the optimization and to take advantage of the complementary properties of the two spaces. Most previous approaches typically focus on either reconstructing in image space [4] or k-space [6]. We evaluate the performance of our method on the recently released fastMRI [20] dataset.

First, we present a brief introduction to parallel MR imaging and review some deep learning methods for parallel MRI

reconstruction in section 2. Next, we provide a description of the GrappaNet model in section 3, followed by a description of our experiments in section 4. Finally, we conclude with a discussion of future work in section 5.

## 2. Background and Related Work

### 2.1. Parallel MRI

MR scanners image a patient’s anatomy by acquiring measurements in the frequency domain using a measuring instrument called a receiver coil. In the MRI literature, these frequency-domain measurements are called *k-space* samples, where  $\mathbf{k}$  refers to the spatial wave number. The image can then be obtained by applying an inverse multidimensional Fourier transform  $\mathcal{F}^{-1}$  to the measured  $k$ -space samples. The underlying image  $\mathbf{x} \in \mathbb{C}^M$  is related to the measured  $k$ -space samples  $\mathbf{k} \in \mathbb{C}^M$  as

$$\mathbf{k} = \mathcal{F}(\mathbf{x}) + \epsilon, \quad (1)$$

where  $\epsilon$  is the measurement noise.

Most modern scanners support parallel imaging: they employ an array of multiple receiver coils that simultaneously obtain  $k$ -space samples from the anatomy being imaged. The  $k$ -space samples measured by each coil are modulated by their sensitivity to the MR signal arising from different regions. In particular, the  $k$ -space sample measured by the  $i$ -th coil is

$$\mathbf{k}_i = \mathcal{F}(S_i \mathbf{x}) + \epsilon_i, i = 1, 2, \dots, N, \quad (2)$$

where  $S_i$  is a complex-valued diagonal matrix encoding the position dependent sensitivity map of the  $i$ -th coil and  $N$  is the number of coils.

Different coils are typically sensitive to different but overlapping regions. It is important to note that the coil sensitivities vary per scan since they depend not only on the configuration of the coils but also on their interaction with the anatomy being imaged.

### 2.2. Accelerated MRI

The speed of MRI acquisition is limited by the number of  $k$ -space samples obtained. This process can be accelerated by obtaining only a subset of the  $k$ -space data:

$$\mathbf{k}_i = M\mathcal{F}(S_i \mathbf{x}) + \epsilon_i, i = 1, 2, \dots, N, \quad (3)$$

where  $M$  is a binary mask operator that selects a subset of the  $k$ -space points. The same mask is used for all coils. Applying an inverse Fourier transform naively to this under-sampled  $k$ -space data results in aliasing artifacts.

Parallel MRI can be used to accelerate imaging by exploiting the redundancies in  $k$ -space samples measured by different coils to estimate the missing  $k$ -space points from the observed points. Various parallel imaging methods have been proposed but they can be divided into two broad classes:

a) *SENSE*-type methods [12] that operate in the image space, and b) *GRAPPA*-type methods [3] that operate locally in  $k$ -space. The latter is relevant to this work.

The GRAPPA algorithm estimates the unobserved  $k$ -space points as a linear combination of the neighboring observed  $k$ -space points from all coils. The same set of weights are used at all spatial locations, which can be seen as a complex-valued convolution in  $k$ -space from  $N$  channels to  $N$  channels, where  $N$  is the number of coils. Formally, the unobserved  $k$ -space points  $\mathbf{k}^u$  are computed from the observed  $k$ -space points  $\mathbf{k}$  by convolving with GRAPPA weights  $G$ :

$$\mathbf{k}^u = G * \mathbf{k}. \quad (4)$$

During acquisition, the central region of  $k$ -space (which corresponds to low spatial frequencies) is fully sampled. This region, called the Auto-Calibration Signal or ACS, is used to estimate the GRAPPA weights  $G$ . We can simulate under-sampling in the ACS by masking out certain  $k$ -space points. Let the simulated observed and unobserved  $k$ -space points in the ACS be  $\mathbf{k}'$  and  $\mathbf{k}^{u'}$  respectively. From equation 4, the convolution of  $G$  and  $\mathbf{k}'$  should be equal to  $\mathbf{k}^{u'}$ . Thus, we can estimate  $G$  by solving the following optimization problem:

$$\hat{G} = \operatorname{argmin}_G \|\mathbf{k}^{u'} - G * \mathbf{k}'\|^2. \quad (5)$$

The knee images in the fastMRI dataset [20] were acquired using machines that employ 15 receiver coils and can generally support  $2\times$  acceleration for imaging of the knee using this approach. Higher acceleration factors lead to aliasing artifacts that cannot be removed by standard parallel imaging methods.

### 2.3. Compressed Sensing for Parallel MRI Reconstruction

Compressed Sensing [2] enables reconstruction of images by using fewer  $k$ -space measurements than is possible with classical signal processing methods by enforcing suitable priors. Compressed sensing has been combined with parallel imaging to achieve higher acceleration factors than those allowed by parallel imaging alone.

Classical compressed sensing methods use sparsity in some transform domain as a prior. Many classical compressed sensing methods operate in the image domain and solve the following optimization problem:

$$\hat{\mathbf{x}} = \operatorname{argmin}_{\mathbf{x}} \frac{1}{2} \sum_i \|\mathcal{F}(S_i \mathbf{x}) - \mathbf{k}_i\|^2 + \lambda \Psi(\mathbf{x}), \quad (6)$$

where  $\Psi$  is a regularization function that enforces a sparsity constraint in some transform domain such as gradients in the image domain. This problem can be solved by iterative gradient descent style methods.

In the last few years, there has been rapid development of deep learning based approaches to MRI reconstruction [4, 9, 11, 15, 18, 21]. MRI reconstruction can be viewed as an inverse problem and several previous research papers have proposed neural networks whose design is inspired by the optimization procedure to solve such an inverse problem [4, 9, 11]. One approach in this direction for the multi-coil reconstruction problem is the Variational Network (VN) [4]. The VN model is a deep neural network, each of whose layers implements one gradient update step for the optimization problem in equation 6. The VN uses pre-computed sensitivity maps and achieves excellent reconstructions at low acceleration factors. Computing sensitivity maps becomes more challenging at higher accelerations, which may limit the maximum acceleration this method can achieve.

An alternate line of work operating in k-space is the RAKI model [1] which replaces the single convolution operation in GRAPPA with a deep convolutional network that is trained independently for each scan. The RAKI method emphasizes the importance of using a scan specific model for multi-coil reconstruction. This method is complementary to our work and can be integrated into the GrappaNet by replacing the GRAPPA layer with the RAKI network.

A comprehensive survey of recent developments in using deep learning for parallel MRI reconstruction can be found in [10].

### 3. GrappaNet

The GrappaNet is a neural network that takes under-sampled, multi-coil k-space data as input and outputs the reconstructed image. Figure 1 shows a diagram of the network architecture that contains three important properties. First, the differentiable GRAPPA layer enables the network to take advantage of the known physical properties of parallel imaging. Next, each convolutional network is applied across all complex-valued views jointly, before being combined in the final stage. This enables the network to take advantage of all the information captured during parallel imaging. Several previous approaches [4, 6], performed reconstruction after collapsing to a single view. Finally, image-to-image mappings using U-Nets are performed in both k-space and image space. Convolutions, pooling, and up-sampling result in very different operations in image space and k-space. We demonstrate in Section 4 that using both these complementary spaces provides improved accuracy.

The network consists of two convolutional neural networks, with the application of the GRAPPA operator in between them. Denoting the input under-sampled k-space data by  $\mathbf{k}$ , the network computes the following function:

$$\mathbf{x} = h \circ f_2(G * f_1(\mathbf{k})), \quad (7)$$

where  $f_1$ , and  $f_2$  are convolutional networks that map multi-coil k-space to multi-coil k-space and  $h$  combines the multi-

coil k-space data to a single image by first applying an inverse fourier transform followed by an root sum-of-squares (RSS) operation (equation 8).

The first network,  $f_1$  takes the multi-coil k-space data with  $R$ -fold under-sampling and maps it to an  $R'$ -fold under-sampled k-space dataset with the same number of coils. The GRAPPA operator,  $G$ , which is separately obtained from the ACS, is then applied to this  $R'$ -fold under-sampled dataset to fill in the missing k-space data. This allows the network to take advantage of the known physical properties of the parallel imaging process.  $R'$  is chosen to be small enough that traditional parallel imaging methods like GRAPPA can reconstruct the image accurately. We use  $R' = 2$  for our experiments.

#### 3.1. U-Net

Both  $f_1$  and  $f_2$  are composed of multiple U-Nets [14], which are convolutional networks that operate at multiple scales. U-Net models and their variants have successfully been used for many image-to-image mapping tasks including MRI reconstruction [5, 7] and image segmentation [13]. The U-Nets used in this work are based on the U-Net baseline models from [20].

A U-Net is useful for image to image mapping tasks like semantic segmentation because the presence of pooling and up-sampling layers allow it to learn useful feature maps at multiple scales and abstraction levels. This multi-resolution feature representation helps the U-Net predict the higher level details of the output at the lowest level of the decoder and gradually adds finer, higher frequency details as the up-sampling layers are applied.

The baseline model, described in [20], used such a U-Net model for MRI image reconstruction. However, that model is only able to perform denoising since it is applied after combining the different views using a root-sum-of-squares (RSS) transform (equation 8). This prevents the baseline model from learning how to combine all of the coils and using the phase information. As a result, the reconstructions from this baseline model are too smooth and lose much of the medically relevant high frequency information (see Figure 3). We show in section 4 that simply applying a U-Net to the real and imaginary data from all coils can significantly improve upon this model. Such a U-Net can potentially learn to combine information from different coils together, which improves performance.

Han et al [6] show that a U-Net can also be applied directly to under-sampled k-space data. Their work was motivated by connections between encoder-decoder models and a classical CS algorithm called the annihilating filter-based low-rank Hankel matrix approach (ALOHA) [8]. The input to the ALOHA U-Net is zero-filled k-space data and the model fills in the missing information. In an approach similar to the fastMRI baseline model [20], Han et al [6]




Figure 1. Left: The full GrappaNet model which takes under-sampled k-space samples as input and outputs the reconstructed image. Right: Details about each of the convolutional networks, which take multi-coil k-space as input and output multi-coil k-space. Here, FT, IFT & RSS refer to 2D Fourier transform, 2D inverse Fourier transform and root sum-of-squares operations (equation 8) respectively.

also apply their U-Net after combining all of the coils into a single coil. Taking insight from algorithms like GRAPPA, we posit that it would be beneficial to apply convolutions directly to the multi-coil k-space data. We show in section 4 that such a model outperforms the baseline models.

The functions  $f_1$  and  $f_2$  apply the following series of operations to the input k-space data (see figure 1): a U-Net in k-space followed by a hard data consistency, inverse 2D Fourier transform to convert to image space, a U-Net in the image space, followed by a 2D Fourier transform and data consistency. Each of the U-Nets map 15 complex-valued channels to 15 complex-valued channels. Here, the hard data consistency operations simply copy all of the observed k-space samples to the correct locations in k-space. This ensures that the model only fills in the missing k-space points.

The function  $h$  combines the reconstructed multi-coil k-space data into a single real-valued image by first applying an inverse 2D Fourier transform to each coil, followed by a root sum-of-squares (RSS) operation. The RSS operation combines all the coils into a single real-valued image:

$$RSS(\mathbf{x}_1, \dots, \mathbf{x}_N) = \left( \sum_{n=1}^N |\mathbf{x}_n|^2 \right)^{1/2}, \quad (8)$$

where  $\mathbf{x}_1, \dots, \mathbf{x}_N$  are the images from the  $N$  coils.

### 3.2. GRAPPA Layer

As explained in the previous section convolutional networks in k-space or image space applied to all coils can, to a limited extent, learn to combine all of the coils. However, as described in 2.2, the coil sensitivities can vary from one imaging examination to another. Traditional parallel imaging methods take this into consideration by estimating distinct sensitivity maps or GRAPPA kernels for each scan. This motivates the need to include a scan-specific component within the neural network that can adapt to differences in the sensitivity profile to improve generalization of the reconstruction model.

We achieve this adaptation by introducing a new neural network layer that we call the *GRAPPA layer*. The GRAPPA layer estimates the GRAPPA kernel from the ACS region and then applies a two dimensional convolution with the estimated kernel. Because the application of GRAPPA is differentiable, the entire network can be trained in an end-to-end fashion using backpropagation.

## 4. Experimental Results

We ran all our experiments on the multi-coil knee MRIs from the fastMRI dataset [20], which consists of raw k-space data from 1594 knee MRI exams from four different MRI machines. The dataset contains two types of MRI sequences that are commonly used for knee exams in clinical practice: a Proton Density (PD) weighted sequence and a Proton Density weighted sequence with Fat Saturation (PDFS). We used the same train, validation and test splits as in the original dataset. The training data consisted of 973 volumes which contained k-space data of different sizes. During training, we omitted k-space data with a width greater than 372, which is about 7% of the training data. We evaluated various models on all test images.

For training our models, we used random masks with  $4\times$  and  $8\times$  accelerations, based on code released with the fastMRI dataset<sup>1</sup>. We experimented with the following models:

1. Classical CS baseline based on Total Variation minimization [20]
2. U-Net baseline model applied to RSS inputs [20]
3. Variational Network model introduced in [4]
4. U-Net applied in k-space to 15 coil input
5. U-Net applied in image space to 15 coil input
6. GrappaNet model

<sup>1</sup><https://github.com/facebookresearch/fastMRI>




Figure 2. Example reconstructions for  $4\times$  under-sampling. The top row shows PD images without fat suppression, and the bottom row shows PDFS images with fat suppression.




Figure 3. Example reconstructions for  $8\times$  under-sampling. The top row shows PD images without fat suppression and the bottom row shows PDFS images with fat suppression.

We used the original implementation of the Variational Network<sup>2</sup>. This code runs the ESPIRiT algorithm [17] to estimate sensitivity maps from the densely sampled ACS region. These maps are used both as input to the network and also to combine the fully sampled coil responses to compute the training targets. For experiments with  $8\times$  accelerations, the input k-space contains very few ACS lines, which yields poor quality sensitivity maps for the Variational

Network. The training targets computed using these poor quality sensitivity maps contain aliasing artifacts that make them unsuitable for training. To mitigate this problem, we always use 30 low frequency lines to compute the training target for  $8\times$  experiments. The sensitivity maps used as inputs to the network are still computed from the ACS region. We did not change the model architecture or training procedure from the original implementation, except for the use of random masks.

<sup>2</sup><https://github.com/VLOGroup/mri-variationalnetwork/>

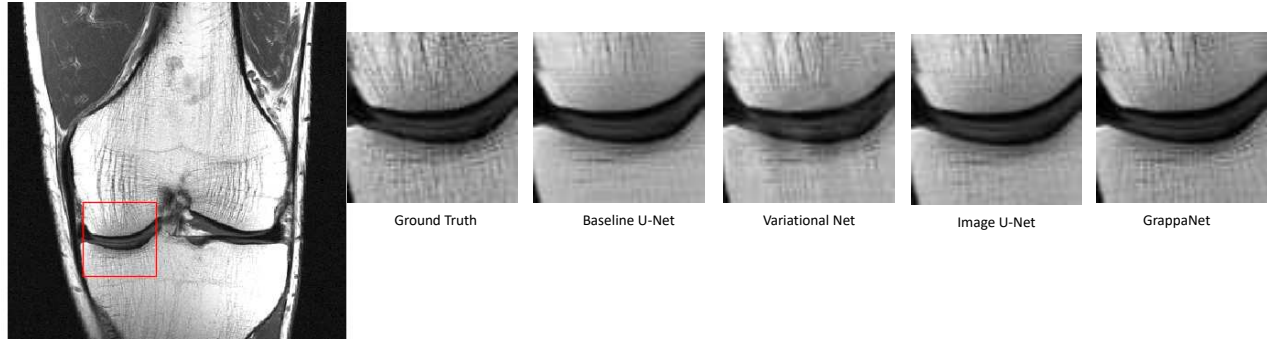


Figure 4. Example reconstructions for  $4\times$  under-sampling with the diagnostically important regions zoomed in.




Figure 5. Left: Ground truth, Middle: Reconstruction from baseline, Right: Reconstruction from GrappaNet. The arrow points to the meniscus region, which appears "filled in" with the baseline method. Highly accurate reconstruction of the meniscus is important for radiologists in diagnosing certain pathologies.

For the k-space U-Net, the image space U-Net, and the GrappaNet models, we followed the training procedure for the baseline models in [20]. To deal with complex-valued inputs, we simply treated the real and imaginary parts as two distinct channels. Hence, 15-coil complex-valued k-space or image data were treated as 30-channel data. These models were trained using the RMSProp [16] algorithm to minimize a linear combination of Structural Similarity (SSIM) [19] and  $L1$  losses:

$$J(\hat{\mathbf{x}}, \mathbf{x}) = -\text{SSIM}(\hat{\mathbf{x}}, \mathbf{x}) + \lambda \|\hat{\mathbf{x}} - \mathbf{x}\|_1, \quad (9)$$

where  $\hat{\mathbf{x}}$  is the reconstruction and  $\mathbf{x}$  is the ground truth image, after cropping to the central  $320 \times 320$  region. Lambda was set to 0.001. The models were trained for 20 epochs with a fixed learning rate of 0.0003. All models were trained on a machine with 8 NVIDIA Volta V100 GPUs using data parallel training for about 3 days.

The U-Net models applied either to 15-coil k-space input or 15-coil image input start with 384 channels, which are doubled after each pooling. The GrappaNet model contains a total of 4 U-Nets, each of which starts with 192 channels. All three models have roughly 480M parameters.

Experimental results are shown in table 1, which lists three metrics that are computed in the same manner as [20]: normalized mean squared error (NMSE), peak signal to noise ratio (PSNR) and structural similarity (SSIM) [19]. All of the proposed models perform significantly better than the baselines. The large difference in performance between a U-Net applied to all 15 coils versus the U-Net baseline

underscores the importance of letting the neural network figure out how to combine the coil images.

The GrappaNet performs best according to all metrics. The improved performance of the GrappaNet can be attributed to the inclusion of the GRAPPA layer to implement parallel imaging within the network.

Some example reconstructions are shown in figures 2 and 3 for  $4\times$  and  $8\times$  accelerations, respectively. Figure 4 and 5 show some of the medically relevant regions zoomed in for  $4\times$  acceleration. The baseline U-Net model is able to remove aliasing artifacts, but this comes at the cost of severe over-smoothing. The reconstruction lacks some of the high frequency detail that is clinically relevant. The reconstructions from the image U-Net model are significantly better than the baseline, but they are not as sharp as the reconstructions from the GrappaNet model.

The Variational Net model makes heavy use of estimated sensitivity maps throughout the network, including in the data consistency terms. It is able to generate good reconstructions with  $4\times$  acceleration, which retains a sufficient number of low frequency lines to estimate sensitivity maps. When the acquisition is accelerated by  $8\times$ , however, the performance degrades significantly since it is not possible to accurately estimate sensitivity maps for this case.

## 5. Conclusion and Future Work

In this paper, we introduced the GrappaNet architecture for multi-coil MRI reconstruction. Multi-coil MRI reconstruction presents an important and challenging problem

Acceleration	Model	NMSE		PSNR		SSIM	
		PD	PDFS	PD	PDFS	PD	PDFS
4-fold	Classical CS baseline	0.0198	0.0951	32.6	27.5	0.693	0.588
	U-Net Baseline	0.0154	0.0525	34.00	29.95	0.815	0.636
	Variational Net	0.0138	0.0262	35.82	33.196	0.919	0.855
	K-Space U-Net	0.0055	0.0114	37.27	36.45	0.927	0.870
	Image U-Net	0.0034	0.0103	39.58	36.97	0.949	0.886
	GrappaNet	<b>0.0026</b>	<b>0.0085</b>	<b>40.74</b>	<b>37.77</b>	<b>0.957</b>	<b>0.891</b>
8-fold	Classical CS baseline	0.0352	0.109	29.6	26.8	0.642	0.551
	U-Net Baseline	0.0261	0.0682	31.5	28.71	0.762	0.559
	Variational Net	0.0211	0.0816	32.12	27.72	0.788	0.675
	K-Space U-Net	0.0189	0.0206	36.45	32.54	0.870	0.807
	Image U-Net	0.0079	0.0160	36.26	34.36	0.886	0.831
	GrappaNet	<b>0.0071</b>	<b>0.0146</b>	<b>36.76</b>	<b>35.04</b>	<b>0.922</b>	<b>0.842</b>

Table 1. Experimental results

due to the prevalence of parallel imaging and the need to make scan-specific adaptations to the neural networks. GrappaNet addresses this challenge by integrating traditional parallel imaging methods with neural networks and training the model end-to-end. This allows the model to generate high fidelity reconstructions even at high acceleration factors.

The GRAPPA kernel used in the GrappaNet model is estimated from the low-frequency lines of k-space and is used as a fixed input to the model. A possible extension to this work could explore methods to optimize the process of estimating the kernel jointly with the rest of the network during training.

Quantitative measures such as NMSE, PSNR, and SSIM only provide an estimate for the quality of the reconstructions. Clinically important details are often subtle and contained in small portions of an MRI. Before techniques such as those presented in this paper can be used in practice, proper clinical validation studies need to be performed to ensure that the use of accelerated MRIs does not degrade the quality of diagnosis.

## References

- [1] Mehmet Akçakaya, Steen Moeller, Sebastian Weingärtner, and Kâmil Uğurbil. Scan-specific robust artificial-neural-networks for k-space interpolation (raki) reconstruction: Database-free deep learning for fast imaging. *Magnetic Resonance in Medicine*, 81(1):439–453, 2019. doi: 10.1002/mrm.27420. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.27420>. 3
- [2] David Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006. 2
- [3] Mark A. Griswold, Peter M. Jakob, Robin M. Heidemann, Mathias Nittka, Vladimir Jellus, Jianmin Wang, Berthold Kiefer, and Axel Haase. Generalized auto-calibrating partially parallel acquisitions (GRAPPA). *Magnetic Resonance in Medicine*, 47(6):1202–1210, 2002. 1, 2
- [4] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P. Recht, Daniel K. Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated MRI data. *Magnetic Resonance in Medicine*, 79(6):3055–3071, 2018. 1, 3, 4
- [5] Yoseob Han and Jong Chul Ye. Framing U-Net via deep convolutional framelets: Application to sparse-view CT. *IEEE Transactions on Medical Imaging*, 37(6), 2018. 3
- [6] Yoseob Han, Jae Jun Yoo, and Jong Chul Ye. Deep learning with domain adaptation for accelerated projection reconstruction MR. *Magnetic Resonance in Medicine*, 80(3), 2018. 1, 3
- [7] Chang Min Hyun, Hwa Pyung Kim, Sung Min Lee, Sungchul Lee, and Jin Keun Seo. Deep learning for undersampled MRI reconstruction. *Physics in medicine and biology*, 63(13), 2018. 3
- [8] K. H. Jin and J. C. Ye. Annihilating filter-based low-rank hankel matrix approach for image inpainting. *IEEE Transactions on Image Processing*, 24(11):3498–3511, 2015. 3
- [9] Kyong Hwan Jin, Michael T. McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional neural network for inverse problems in imaging. *CoRR*, abs/1611.03679, 2016. URL <http://arxiv.org/abs/1611.03679>. 3

- [10] Florian Knoll, Kerstin Hammernik, Chi Zhang, S. Möller, Thomas Pock, Daniel K. Sodickson, and Mehmet Akçakaya. Deep learning methods for parallel magnetic resonance image reconstruction. *CoRR*, abs/1904.01112, 2019. 3
- [11] Dong Liang, Jing Cheng, Ziwen Ke, and Leslie Ying. Deep mri reconstruction: Unrolled optimization algorithms meet neural networks. *arXiv preprint arXiv:1907.11711*, 2019. 3
- [12] Klaas P Pruessmann, Markus Weiger, Markus B Scheidegger, and Peter Boesiger. SENSE: sensitivity encoding for fast MRI. *Magnetic resonance in medicine*, 42(5), 1999. 1, 2
- [13] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention*, 2015. 3
- [14] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI 2015: Medical Image Computing and Computer-Assisted Intervention*, volume 9351 of *Lecture Notes in Computer Science*, pages 234–241. Springer, 2015. 3
- [15] Jo Schlemper, Jose Caballero, Joseph V. Hajnal, Anthony N. Price, and Daniel Rueckert. A deep cascade of convolutional neural networks for MR image reconstruction. *Information Processing in Medical Imaging*, 2017. 1, 3
- [16] Tijmen Tieleman and Geoffrey Hinton. Lecture 6.5 - rmsprop: Divide the gradient by a running average of its recent magnitude. COURSERA: Neural networks for machine learning, 2012. 6
- [17] Martin Uecker, Peng Lai, Mark J Murphy, Patrick Virtue, Michael Elad, John M Pauly, Shreyas S Vasanawala, and Michael Lustig. ESPIRiT -an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. *Magnetic resonance in medicine*, 71(3), 2014. 5
- [18] Shanshan Wang, Zhenghang Su, Leslie Ying, Xi Peng, Shun Zhu, Feng Liang, Dagan Feng, and Dong Liang. Accelerating magnetic resonance imaging via deep learning. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2016. 1, 3
- [19] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. 6
- [20] Jure Zbontar, Florian Knoll, Anuroop Sriram, Matthew J. Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdal, Adriana Romero, Michael Rabbat, Pascal Vincent, James Pinkerton, Duo Wang, Nafissa Yakubova, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastmri: An open dataset and benchmarks for accelerated MRI. *CoRR*, abs/1811.08839, 2018. URL <http://arxiv.org/abs/1811.08839>. 1, 2, 3, 4, 6
- [21] Bo Zhu, Jeremiah Z. Liu, Stephen F. Cauley, Bruce R. Rosen, and Matthew S. Rosen. Image reconstruction by domain-transform manifold learning. *Nature*, 555(7697), 2018. 1, 3