

A Hybrid Frequency-domain/Image-domain Deep Network for Magnetic Resonance Image Reconstruction

Roberto Souza^{1,2} and Richard Frayne^{1,2}

¹Radiology and Clinical Neuroscience, Hotchkiss Brain Institute, University of Calgary, Calgary, AB, Canada

²Seaman Family MR Research Centre, Foothills Medical Centre, Calgary, AB, Canada

Email: {roberto.medeirosdeso,rfrayne}@ucalgary.ca

Abstract—Decreasing magnetic resonance (MR) image acquisition times can potentially make MR examinations more accessible. Compressed sensing (CS)-based image reconstruction methods decrease MR acquisition time by reconstructing high-quality images from data that were originally sampled at rates inferior to the Nyquist-Shannon sampling theorem. Deep-learning methods have been used to solve the CS MR reconstruction problem. These proposed methods are able to quickly reconstruct images in a single pass using an appropriately trained network. A variety of different network architectures (*e.g.*, U-nets and Residual U-nets) have been proposed to tackle the CS reconstruction problem. A drawback of these architectures is that they typically only work on image domain data. For undersampled data, the images computed by applying the inverse Fast Fourier Transform (iFFT) are aliased. In this work we propose a hybrid architecture, termed W-net, that works both in the k-space (or frequency-domain) and the image (or spatial) domains. Our network is composed of a complex-valued residual U-net in the k-space domain, an iFFT operation, and a real-valued U-net in the image domain. Our experiments demonstrated, using MR raw k-space data, that the proposed hybrid approach can potentially improve CS reconstruction compared to deep-learning networks that operate only in the image domain. In this study we compare our method with four previously published deep neural networks and examine their ability to reconstruct images that are subsequently used to generate regional volume estimates. Our technique was ranked second in the quantitative analysis, but qualitative analysis indicated that our reconstruction performed the best in hard to reconstruct regions, such as the cerebellum. All images reconstructed with our method were successfully post-processed, and showed good volumetry agreement compared with the fully sampled reconstruction measures.

I. INTRODUCTION

Magnetic resonance (MR) is a key medical imaging modality that has critical roles in both patient care and medical research. MR scanner installations, however, are expensive. Another major limitation to MR imaging is the comparatively long image acquisition times, especially when compared to other modalities like computerized tomography. Lengthy acquisition times make MR less patient friendly and increase the per patient examination cost. MR-based compressed sensing (CS) methods seek to leverage the implicit sparsity of medical images [1], potentially allowing for significant k-space undersampling during acquisition, and by consequence, reducing examination times.

Traditional MR CS reconstruction techniques are iterative algorithms that usually require a sparsifying transform that when combined with regularization parameters are able to find a solution for these ill-posed inverse problems [1], [2]. The drawback of these iterative approaches, however, is that they are time-consuming, making them more difficult to incorporate in a near real-time MR imaging scenario (*i.e.*, where images are reconstructed and displayed on the scanner during the procedure). Deep learning [3] is considered a new frontier in image reconstruction [4]. It has the advantage of being able to rapidly reconstruct images in a single-pass using a suitably trained network. Some deep-learning based reconstruction methods have arguably surpassed traditional iterative CS reconstruction techniques [5].

A few different deep learning approaches have been recently proposed to tackle the CS reconstruction problem. Jin *et al.* [6] proposed to use a U-net [7]. Lee *et al.* experimentally showed that residual U-nets can potentially improve image reconstruction. Residual blocks have subsequently been incorporated in the majority of the latest studies (*cf.*, [5], [8], [9]). Yang *et al.* [5] proposed a deep de-aliasing generative adversarial network (DAGAN) that uses a residual architecture as the network generator responsible for reconstructing the images associated to a loss function that has four components: an image domain loss, a frequency domain loss, a perceptual loss, and an adversarial loss. Quan *et al.* [8] proposed a generative adversarial network (GAN) with a cyclic loss [10]. Their method consists of a cascade of a reconstruction network followed by a refinement network with a cyclic loss component that tries to enforce that the model is bijective. Schlemper *et al.* [9] proposed a deep cascade of convolutional neural networks (CNNs) that has data consistency (DC) blocks between consecutive subnetworks in the cascade. The DC blocks potentially reduce the issue of overfitting, therefore allowing the training of deeper models. The aforementioned techniques principally work in the image domain, with a few exceptions, where k-space domain information is used in the loss function and/or to implement DC layers. All of these networks [5], [8], [9] receive as input the undersampled k-space zero-filled reconstruction, and output an unaliased image.

Zhu *et al.* [11] recently proposed a method that tries to learn the domain transform. Their method first processes the

undersampled input data in k-space, learns the inverse discrete Fourier transform (iDFT), and then refines the result in the image domain. In the case of CS reconstruction, the domain transform can be considered as learning an approximation for the iDFT for undersampled k-space data. The domain transform is modeled as a sequence of connected layers, and the image domain refinement is modeled as a series of convolutional layers. A disadvantage of this approach is that the domain transform has a quadratic complexity with respect to the size of the input image. For example, when dealing with 256×256 images, the number of learned parameters in their model would be greater $> 10^{10}$. The quadratic order of the algorithm makes it challenging to use their model for typical MR image sizes due to current hardware limitations. Recent work [12], [13] has tried to decompose the 2-dimensional transform into two 1-dimensional transforms, therefore mitigating the parameter complexity problem.

Eo *et al.* [14] developed a hybrid architecture named KIKI-net that cascades flat unrolled sub-networks that alternate between k-space and image domains. A similar approach has also been used for computed tomography reconstruction [15]. In a further investigation of KIKI-net [16], other possible domain configurations for the sub-networks in the cascade were investigated and their results indicated that starting the cascade with an image-domain sub-network may be advantageous.

Based on these studies, we hypothesize that a hybrid approach that works with both information as presented in k-space domain and image domain can improve MR CS reconstruction. In this work, we propose a model, termed hybrid W-net, that consists of a cascade of a k-space domain network, which is implemented through a residual U-net, and an image domain network, implemented through a U-net, connected by the magnitude of the iDFT operation. Our model does not need to learn the domain transform, which essentially reduces our model parameter complexity by a factor of 10^3 compared to [11]. Our proposed method takes advantage of information as presented in k-space and image domain, as opposed to other image domain only approaches [3], [5], [6], [8], [9]. K-Space and image domain information are equivalent, because they are related by a global linear transformation. Operating in both domains with non-linear methods, however, can potentially be advantageous and improve the network learning. Another potential practical advantage is that our hybrid formulation is compatible with MR parallel imaging approaches due the fact that multi-channel coil sensitivity introduces data correlation and every data sample may be represented by the convolution of the neighboring samples, which can potentially be learned by a k-space CNN [17].

In this work, we compare our proposal against four recently published, publicly available, deep learning-based reconstruction methods [5], [6], [8], [9]. These approaches were evaluated at 75% and 80% undersampling levels (corresponding to acquisition acceleration factors of $4\times$ and $5\times$). The experiments were done using MR raw-data acquired from subjects scanned using a volumetric $T1$ -weighted MR imaging sequence. The proposed reconstruction code and the

dataset used in the experiments are publicly available at <https://github.com/rmsouza01/Hybrid-CS-Model-MRI> and <https://sites.google.com/view/calgary-campinas-dataset> [18], respectively.

II. HYBRID NETWORK

The flowchart of our proposed method is depicted in Figure 1. There are three main components to our approach: 1) a k-space (or frequency-domain) network that essentially tries to fill missing k-space samples, 2) the magnitude of the iDFT, and 3) an image-domain network that acts as an anti-aliasing filter. Topographically our implementation consists of two U-Nets [7] connected by the iDFT; we also describe this arrangement as a W-Net (Figure 2). These components along with the proposed network loss function are described in the following subsections.

A. Frequency-domain Network

The frequency-domain network, f_{cnn1} , attempts to recover a fully sampled k-space, $\hat{F}_{norm}(k_x, k_y)$, given undersampled k-space data, F_u . This can be mathematically formalized as:

$$\hat{F}_{norm}(k_x, k_y) = f_{cnn1}[F_{u_{norm}}], \quad (1)$$

where $F_{u_{norm}}$ is the normalized undersampled k-space data given by:

$$F_{u_{norm}} = \frac{F_u - \mu_{F_{u_{train}}}}{\sigma_{F_{u_{train}}}}, \quad (2)$$

where $\mu_{F_{u_{train}}}$ and $\sigma_{F_{u_{train}}}$ are the average and standard-deviation of undersampled k-spaces in the training set. The specific architecture used for the frequency-domain network is a residual U-net (Figure 2, left side). The input complex k-space image is split in two-channels: one for the real and other for the imaginary components of the k-space data.

B. Magnitude of the iDFT

Before applying the iDFT, we have to undo the previous k-space normalization step. Adding a constant to the k-space data results in superposition of an impulse, $\delta(\cdot)$, signal to the image after transformation. Undoing the normalization is accomplished by:

$$\hat{F}(k_x, k_y) = F_{u_{norm}} \times \sigma_{F_{u_{train}}} + \mu_{F_{u_{train}}}. \quad (3)$$

The next step is to transform from the frequency domain to the image domain using the iDFT and magnitude operations:

$$\hat{f}_0 = \|\mathcal{F}^{-1}(\hat{F})\| \quad (4)$$

where \mathcal{F}^{-1} represents the iDFT operation and \hat{f}_0 is the initial estimate of the reconstructed image.

This component of our model has no trainable parameters, and the iDFT is efficiently computed using the fast Fourier transform (FFT) algorithm running on a graphics processing unit (GPU).




Fig. 1. Flowchart of the proposed methodology. The frequency domain network can be seen as k-space interpolation to fill the missing values. The image domain network acts as an anti-aliasing filter to further improve the image reconstruction obtained from the first network.

C. Image Domain Network

The last component of our method is the image domain network (f_{cnn2}). In order to improve training convergence of the network, we again normalize the initial estimate of the reconstructed image obtained in the previous step:

$$\hat{f}_{0_{norm}} = \frac{\hat{f}_0 - \mu_{f_{0_{train}}}}{\sigma_{f_{0_{train}}}}. \quad (5)$$

where $\mu_{f_{0_{train}}}$ and $\sigma_{f_{0_{train}}}$ are the mean and standard-deviation of the reconstructed images in the training set. The normalized image $\hat{f}_{0_{norm}}$ is fed as input to the image domain network:

$$\hat{f}(x, y) = f_{cnn2}[\hat{f}_{0_{norm}}]. \quad (6)$$

The final reconstructed image is $\hat{f}(x, y)$. The architecture used for the image domain network is a U-net (Figure 2, right side).

D. Loss Function

Our loss function was a weighted sum of normalized root mean squared errors (NRMSE) in each domain given by:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N w \times NRMSE(F_i, \hat{F}_i) + (1-w) \times NRMSE(f_i, \hat{f}_i), \quad (7)$$

where F_i and f_i are the reference fully sampled k-space and image reconstruction, respectively, of the i -th sample in the training set, and N is the number of training samples. In our experiments $w = 0.001$. These values were chosen to ensure greater weight is applied to the final network output.

III. EXPERIMENTAL SETUP

A. Network Implementations

We compared our method, which we will refer to simply as Hybrid method, against 1) a plain vanilla U-net [7] with a residual connection, referred to as UNET; 2) RefineGAN [8]; 3) DAGAN [5]; and 4) Deep-Cascade [9] with a cascade of five CNNs and five DC blocks, which is the network configuration that the authors reported best results.

We used the Keras application program interface [19] using TensorFlow as backend [20] to implement our hybrid network and the UNET. For RefineGAN, DAGAN and Deep-Cascade, we used the source code provided by the authors. All networks were trained using our data for acceleration factors of $4\times$ and

$5\times$ using two-dimensional Gaussian undersampling patterns. The networks were trained and tested on Amazon Elastic Compute Cloud services using a p3.2xlarge¹ instance, which has a NVIDIA Tesla V100 GPU.

B. Training, Validation and Testing Dataset

Our dataset consists of 45 volumetric $T1$ -weighted, fully sampled k-space datasets acquired on a clinical MR scanner (Discovery MR750; General Electric (GE) Healthcare, Waukesha, WI) that were collected as part of the ongoing Calgary Normative Study [21]. The data was acquired with a 12-channel imaging coil and an acquisition matrix of size $256 \times 218 \times 170$. Data were zero-filled to an image matrix size of $256 \times 256 \times 170$. The multi-coil k-space data was reconstructed using vendor supplied tools (Orchestra Toolbox; GE Healthcare). Coil sensitivity maps were normalized to produce a single complex-valued image set that could be back-transformed to regenerate complex k-space samples. In our experiments, we performed retrospective undersampling, effectively simulating a single-coil acquisition scenario. Our train/validation/test data split was 25/10/10, equivalent to 4,524 slices/1,700 slices/1,700 slices. There was no overlap of same subject slices in the train, validations and test sets.

C. Performance Metrics

The performance metrics used to assess the networks were:

- NRMSE

$$NRMSE(\hat{f}, f) = \frac{\sqrt{\frac{1}{M} \sum_{i=1}^M [\hat{f}(i) - f(i)]^2}}{\max(f) - \min(f)}, \quad (8)$$

where M is the number of pixels (or voxels) in the image.

- Structural Similarity (SSIM) [22]

$$SSIM(\hat{f}, f) = \frac{(2\mu_f \mu_{\hat{f}} + c_1)(2\sigma_{f\hat{f}} + c_2)}{(\mu_f^2 + \mu_{\hat{f}}^2 + c_1)(\sigma_f^2 + \sigma_{\hat{f}}^2 + c_2)}, \quad (9)$$

where c_1 and c_2 are two variables used to stabilize the division and μ and σ^2 represent mean and variance values of the gray-level intensities of the images.

- Peak Signal to Noise Ratio (PSNR):

$$PSNR(\hat{f}, f) = 20 \log_{10} \left(\frac{\max(f)}{\sqrt{\frac{1}{M} \sum_{i=1}^M [\hat{f}(i) - f(i)]^2}} \right). \quad (10)$$

¹<https://aws.amazon.com/ec2/instance-types/p3/>




Fig. 2. Architecture of our hybrid model. The k-space network uses a 5×5 convolution kernel size and the image domain network uses a 3×3 kernel. Due to its topography consisting of two U-Nets, we term this hybrid method a W-Net.

These metrics were used because they are commonly used to assess image reconstruction. The higher the SSIM and PSNR values the better the result. For NRMSE, smaller values represent better reconstructions.

D. Volumetric Assessment

For the top two performing reconstruction techniques, we performed volumetric analysis with a commonly used software for neuroimaging analysis (FreeSurfer [23]). We considered the fully sampled reconstruction results as our gold-standard. Only the ten test volumes were analyzed. We recorded number of processing failures and analyzed the average absolute deviation of total intra-cranial, white-matter, gray-matter, hippocampus, and amygdala volumes. This analysis is similar to the semantic interpretability score proposed in [24].

E. Statistical Methods

We report mean and standard deviation of all average measures. We used a one-way analysis of variance (ANOVA) to determine statistically significant changes and *post-hoc* paired *t*-tests to determine statistically significant pair-wise differences. A p -value < 0.01 was used to determine statistical significance.

IV. RESULTS

A. Quantitative Assessment

The reconstruction metrics are summarized in Table I. Deep-Cascade achieved the best quantitative results followed by our Hybrid approach across all performance metrics. The one-way ANOVA tests showed statistically significant differences ($p < 0.01$) across all metrics and acceleration factors. The paired *t*-tests showed that Deep-Cascade was statistically better




Fig. 3. Average NRMSE distribution along the slices. Edge slices that do not have much signal have higher average errors.

($p < 0.01$) than the other methods in the comparison and our Hybrid, while worse than Deep-Cascade, was statistically better than UNET, DAGAN, and RefineGAN ($p < 0.01$). We note that the absolute differences between Deep-Cascade and Hybrid networks, while statistically significant, were small.

The average NRMSE across the slices in the ten test volumes is depicted in Figure 3. Over the central slices NRMSE followed a predictable pattern. Towards the edges, where little of the brain was present, the error increased. Sample reconstructions for each technique are shown in Figure 4. Qualitatively the Hybrid network produces the most pleasing image, particularly in regions with large signal differences like the cerebellum. This finding was consistent across all reconstructions in the test set.




Fig. 4. Sample reconstructions with a special highlight on the cerebellum region, where differences are more noticeable.

B. Volumetric Analysis

Volumetric analysis results are summarized in Table II. An example of a subject where processing failed for the fully sampled reconstruction and the Deep-Cascade network (for both $4\times$ and $5\times$ acceleration factors) is presented in Figure 5. Box-plots of the estimated volumes of total intra-cranial volume, white-matter, gray-matter, hippocampus, and amygdala volumes are shown in Figure 6. These plots represent the distribution of eight volumes from only eight subjects in the test set (because volumetric analysis failed to process two fully sampled image sets). Due to the reduced sample size ($n = 8$) and because of reduced statistical power, we did not perform statistical tests on the volumetric data.

V. DISCUSSION

Our hybrid method operates in frequency domain using a residual U-net and image domain using a U-net. The two networks are connected through the magnitude of the iDFT operation, and the model is fully trained end-to-end. An example for one subject of the input, intermediary and output results of our model are depicted in Figure 7. Improvement in the image quality and NRMSE is noticeable at the output of the frequency domain network ($NRMSE = 1.9\%$), and further improvement can be seen at the output of the image domain network ($NRMSE = 1.6\%$). NRMSE for the input image was 3.2% .

TABLE I

SUMMARY OF THE RESULTS FOR THE DIFFERENT ARCHITECTURES AND DIFFERENT ACCELERATION FACTORS. THE TOP TWO RESULTS FOR EACH METRIC AND ACCELERATION FACTOR ARE EMBOLDENED.

Model	SSIM	NRMSE (%)	PSNR
Speed-up: $4\times$			
UNET	0.977 ± 0.062	2.33 ± 1.04	33.28 ± 3.14
DAGAN	0.963 ± 0.105	2.93 ± 1.47	31.33 ± 3.11
RefineGAN	0.979 ± 0.068	1.90 ± 1.30	35.44 ± 3.71
Deep-Cascade	0.986 ± 0.054	1.20 ± 1.06	39.51 ± 3.35
Hybrid	0.985 ± 0.055	1.36 ± 1.01	38.23 ± 3.32
Speed-up: $5\times$			
UNET	0.966 ± 0.096	2.73 ± 1.17	31.88 ± 3.13
DAGAN	0.949 ± 0.110	3.87 ± 1.44	28.69 ± 2.66
RefineGAN	0.973 ± 0.082	2.27 ± 1.40	33.84 ± 3.83
Deep-Cascade	0.982 ± 0.068	1.45 ± 1.11	37.67 ± 3.20
Hybrid	0.981 ± 0.060	1.65 ± 1.09	36.50 ± 3.23

TABLE II

SUMMARY OF VOLUME ANALYSIS PROCESSING FAILURES.

Model	Acceleration factor	# failures (%)
Fully sampled	$1\times$	2 (20%)
DEEP-CASCADE	$4\times$	2 (20%)
	$5\times$	1 (10%)
Hybrid	$4\times$	0 (0%)
	$5\times$	0 (0%)

Our Hybrid model achieved the second best metrics in the quantitative assessment, losing only to Deep-Cascade.




Fig. 5. Sample result of volumetric analysis failures in one subject. Processing failed for the fully sampled reconstruction, and Deep-Cascade (acceleration factors of 4× and 5×). * represents an analysis that failed.




Fig. 6. Summary of volume analysis processing results.

RefineGAN was the third best method. It is important to mention that in the original RefineGAN paper the authors did not use a test set. They reported results on a validation set. The poorest performing technique was DAGAN, though it still had $NRMSE < 4\%$ and $SSIM > 0.94$. It is also important to highlight that in the original DAGAN paper, the authors did not use MR raw k-space data in their experiments. They computed the FFT of magnitude images followed by retrospective undersampling. This is not a realistic scenario, because when applying the FFT operator to magnitude images, the output is a k-space with Hermitian symmetry, while raw k-space is not Hermitian.

Visual assessment of the reconstructions (Figure 4) indicate that Hybrid and Deep-Cascade reconstruction are the best at

preserving fine details as can be more noticeably seen in the cerebellum region. With the Hybrid reconstruction edges seem smoother, which might be an explanation to the fact that volumetric analysis software was able to successfully process all ten image dates sets, while it failed twice when processing the fully sampled reconstruction, and three times with Deep-Cascade. A similar finding was found in [24], and can be explained by the fact that the smooth appearance of the Hybrid reconstructions can be advantageous for the segmentation task. The quality of the brain structure segmentations vary according to the reconstruction. This is specially noticeable in the brain extraction step (Figure 5).

The deviation of the volumes estimated from the fully sampled reconstruction measurements were used as gold-standard




Fig. 7. Illustration of the images and their corresponding k-spaces at different stages of the Hybrid network. From the left to the right: Input undersampled k-space reconstruction ($NRMSE = 3.2\%$), result of the frequency domain U-net ($NRMSE = 1.9\%$), result of the image domain U-net ($NRMSE = 1.6\%$), and fully sampled reference reconstruction.

(reference) measure for the eight successful subjects. Good agreement was found between the reference and volumetric results for Deep-Cascade and Hybrid image sets. For total intra-cranial volume the average absolute volume deviation was $< 0.1\%$ in all subjects. For cerebral white-matter the absolute volume deviation was $< 6\%$ in all subjects, and for gray-matter, it was $< 4\%$ in all subjects. The average differences for the hippocampus and amygdala volumes was between 2% and 3%. These results are a suggestive of the feasibility of incorporating a $5\times$ acceleration factor in a clinical or research imaging setting.

In terms of processing time, we did not perform a systematic comparison due to the fact that the different methods were implemented in different deep learning frameworks and many of the networks were not optimized. Nevertheless, all five techniques assessed take only a few seconds to reconstruct an entire MR volume using a cloud-based GPU (Amazon Elastic Compute Cloud). In terms of processing times, UNET, DAGAN, Deep-Cascade and Hybrid took < 6 hours to train, while RefineGAN training time was ≈ 72 hours.

VI. CONCLUSION

In this work, we proposed a hybrid frequency domain/image domain CS MR reconstruction method that leverages the information of the iDFT mathematical formulation, essentially reducing the number of parameters of our model by orders of magnitude (for a 256×256 image, the Hybrid W-net model required $10^3 \times$ fewer coefficients compared to [11]). Our model method was the second best in the quantitative comparison and it was the only one that did not fail in the

volumetric analysis processing pipeline. Also, our Hybrid W-net model produced the most visually pleasing images (Figure 4).

As future work, we would like to investigate the use of DC layers in our architecture and add an adversarial component for potentially improving reconstruction of high-frequency contents of the image. We will also extend our model to a more generic framework that can potentially deal with the parallel imaging scenario (*cf.*, [25], [26]).

ACKNOWLEDGMENTS

The authors would like to thank the Natural Science and Engineering Council of Canada (NSERC) for operating support, NVidia for providing a Titan V GPU, and Amazon Web Services for access to cloud-based GPU services. We would also like to thank Dr. Louis Lauzon for setting up the script to save the raw MR data at the Seaman Family MR Centre. R.S. was supported by an NSERC CREATE I3T Award and currently holds the T. Chen Fong Fellowship in Medical Imaging from the University of Calgary. R.F. holds the Hopewell Professorship of Brain Imaging at the University of Calgary.

REFERENCES

- [1] 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.
- [2] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 72–82, 2008.
- [3] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, p. 436, 2015.

- [4] G. Wang, J. C. Ye, K. Mueller, and J. A. Fessler, "Image reconstruction is a new frontier of machine learning," *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1289–1296, 2018.
- [5] G. Yang, S. Yu, H. Dong, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, and Y. Guo, "DAGAN: Deep de-aliasing generative adversarial networks for fast compressed sensing MRI reconstruction," *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1310–1321, 2018.
- [6] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Transactions on Image Processing*, vol. 26, no. 9, pp. 4509–4522, 2017.
- [7] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [8] T. M. Quan, T. Nguyen-Duc, and W.-K. Jeong, "Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss," *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1488–1497, 2018.
- [9] J. Schlemper, J. Caballero, J. V. Hajnal, A. N. Price, and D. Rueckert, "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE transactions on Medical Imaging*, vol. 37, no. 2, pp. 491–503, 2018.
- [10] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," *arXiv preprint*, 2017.
- [11] B. Zhu, J. Z. Liu, S. F. Cauley, B. R. Rosen, and M. S. Rosen, "Image reconstruction by domain-transform manifold learning," *Nature*, vol. 555, no. 7697, p. 487, 2018.
- [12] T. Eo, H. Shin, T. Kim, Y. Jun, and D. Hwang, "Translation of 1d inverse fourier transform of k-space to an image based on deep learning for accelerating magnetic resonance imaging," in *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer, 2018, Conference Proceedings, pp. 241–249.
- [13] J. Schlemper, I. Oksuz, J. Clough, J. Duan, A. King, J. Schanbel, J. Hajnal, and D. Rueckert, "dAUTOMAP: Decomposing AUTOMAP to achieve scalability and enhance performance," in *International Society for Magnetic Resonance in Medicine (ISMRM)*, 2019.
- [14] T. Eo, Y. Jun, T. Kim, J. Jang, H.-J. Lee, and D. Hwang, "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.
- [15] J. Adler and O. Öktem, "Learned primal-dual reconstruction," *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1322–1332, 2018.
- [16] R. Souza, R. M. Lebel, and R. Frayne, "A hybrid, dual domain, cascade of convolutional neural networks for magnetic resonance image reconstruction," in *Proceedings of The 2nd International Conference on Medical Imaging with Deep Learning*, ser. Proceedings of Machine Learning Research, vol. 102. London, United Kingdom: PMLR, 08–10 Jul 2019, pp. 437–446.
- [17] P. Zhang, F. Wang, W. Xu, and Y. Li, "Multi-channel generative adversarial network for parallel magnetic resonance image reconstruction in k-space," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, Conference Proceedings, pp. 180–188.
- [18] R. Souza, O. Lucena, J. Garrafa, D. Gobbi, M. Saluzzi, S. Appenzeller, L. Rittner, R. Frayne, and R. Lotufo, "An open, multi-vendor, multi-field-strength brain MR dataset and analysis of publicly available skull stripping methods agreement," *NeuroImage*, 2017.
- [19] F. Chollet *et al.*, "Keras," <https://github.com/fchollet/keras>, 2015.
- [20] M. Abadi *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015.
- [21] A. Tsang, C. A. Lebel, S. L. Bray, B. G. Goodyear, M. Hafeez, R. C. Sotero, C. R. McCreary, and R. Frayne, "White matter structural connectivity is not correlated to cortical resting-state functional connectivity over the healthy adult lifespan," *Frontiers in Aging Neuroscience*, vol. 9, p. 144, 2017.
- [22] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [23] B. Fischl, "Freesurfer," *NeuroImage*, vol. 62, no. 2, pp. 774–781, 2012.
- [24] M. Seitzer, G. Yang, J. Schlemper, O. Oktay, T. Würfl, V. Christlein, T. Wong, R. Mohiaddin, D. Firmin, J. Keegan *et al.*, "Adversarial and perceptual refinement for compressed sensing mri reconstruction," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2018, pp. 232–240.
- [25] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: sensitivity encoding for fast MRI," *Magnetic Resonance in Medicine*, vol. 42, no. 5, pp. 952–962, 1999.
- [26] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, and A. Haase, "Generalized autocalibrating partially parallel acquisitions (GRAPPA)," *Magnetic Resonance in Medicine*, vol. 47, no. 6, pp. 1202–1210, 2002.