K-space refinement in deep learning MR reconstruction via regularizing scan specific SPIRiT-based self consistency

Kanghyun Ryu[1], Cagan Alkan[2], Chanyeol Choi[3], Ikbeom Jang[4] , Shreyas Vasanawala[1]

[1] Department of Radiology, Stanford University
[2] Department of Electrical Engineering, Stanford University
[3] Electrical Engineering and Computer Science, Massachusetts Institute of Technology (MIT)
[4] Massachusetts General Hospital (MGH) / Harvard Medical School

kanghyun@stanford.edu, calkan@stanford.edu, chanyeol@mit.edu, ikbeom.jang@mgh.harvard.edu
vasanawala@stanford.edu

Abstract

Deep Learning (DL) based reconstruction using unrolled neural networks has shown great potential in accelerating magnetic resonance imaging (MRI). However, one of the major drawbacks is the loss of high-frequency details and textures in the output. In this paper, we propose a novel refinement method based on SPIRiT (Iterative Self-consistent Parallel Imaging Reconstruction from Arbitrary k-Space) formulation to reduce the k-space errors and enable reconstruction of improved high-frequency image details and textures. The proposed scheme constrains the DL output to satisfy the neighborhood relationship in the frequency space (k-space) which can be easily calibrated in the auto-calibration (ACS) lines, and corrects the underestimation in the peripheral region of the k-space as well as reduce structured k-space errors. We show that our method enables the reconstruction of sharper images with significantly improved high-frequency components measured by HFEN and GMSD while maintaining overall error in the image measured by PSNR and SSIM.

1. Introduction

Following the success of deep learning (DL) in a wide range of applications, neural network-based machine-learning techniques have received keen interest as a means of accelerating magnetic resonance imaging (MRI) [1, 2, 3, 4, 5, 6, 7, 8]. Among various network architectures, unrolled neural networks have shown superior efficacy as demonstrated in large-scale MRI reconstruction challenges [9, 10]. Unrolled neural networks enabled significant improvement with ideas borrowed from compressed sensing (CS) reconstruction in the form of unrolled iterative optimization [7, 8, 11], which we will refer to as the DL-based method for abbreviation. For details on the developments of unrolled neural networks, refer to suggested survey papers [12, 13].

While the DL-based methods showed significantly improved image quality compared to the traditional CS methods, one of the major drawbacks of these methods is the over-smoothing behavior and the loss of high-frequency details [9, 10, 14]. Global or quantitative measures such as MSE and SSIM are commonly used as loss functions for training the networks. However, they are known to represent the human perceptual system poorly and tend to be oblivious of important local features such as subtle details, often resulting in blurry images [5, 11, 14].

There have been several studies that explored feature losses with generative adversarial networks (GANs), which are notable for generating high-quality realistic images [4, 5]. However, GAN-based methods are shown to have the possibility of creating hallucinations of artificial structures [15] and increase ambiguity in the reconstructed results [16]. Perceptual losses, which use pre-trained networks on a large data-set to compute the feature loss [5, 17] have recently gained interest in this context.

Recently, DL-based methods that operate both in k-space and image domains (also called dual-domain methods) have gained interest in single-coil MR reconstruction. These methods are based on the intuition that aliasing artifacts generated in the image domain are structural and non-local, and sole image domain restoration may be insufficient [3, 18, 19]. However, these approaches require more parameters and higher computation load during training. Thus, it is also an open question how well these dual-domain approaches may perform for multi-coil MR reconstruction or how well they can generalize to anatomical changes and

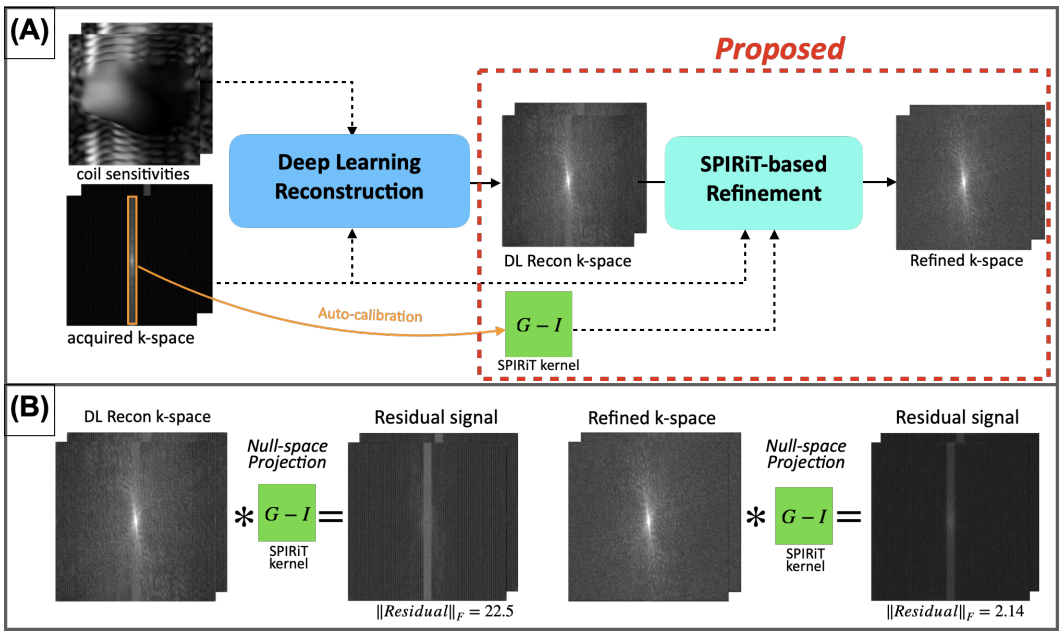

Figure 1. The overall scheme of the current method. (A) The proposed SPIRiT-based refinement is applied on DL-reconstructed k-space to improve and refine. The additional input to the refinement is the acquired k-space and the SPIRiT kernel ($G - I$), which is auto-calibrated in the center lines of the k-space. (B) The underlying principle of the method is to reduce the residual signal when convolving the k-space with the SPIRiT kernel. This can constrain the resulting k-space by enforcing self-consistency (or neighborhood consistency). As can be seen, our method greatly reduces the Frobenius norm of the residual signal.

variations in training and test domain at a large scale, due to a high number of parameters.

Traditionally, in the context of multi-coil k-space restoration, GRAPPA [20] is an auto-calibrating technique based on the local k-space kernels, which utilizes the learned correlation between multiple channels in local areas and fill the missing k-space values by a linear combination of the acquired local data from multiple coils. Iterative self-consistent parallel imaging reconstruction (SPIRiT) [21] is based on the GRAPPA framework but formulated as an inverse problem that can reconstruct data from arbitrary k-space trajectories. Extended versions of SPIRiT such as ACS-LORAKS [22] and ESPIRiT [23] have been presented, utilizing more structured matrices and eigenvalue decomposition using the information of limited spatial support or slowly varying phase.

This work aims to introduce a novel k-space refinement approach based on self neighborhood consistency from SPIRiT formulation that reduces the structured k-space error from the DL-based method. The main advantage of this method is that it can significantly enhance the texture and high-frequency image features without using sophisticated GAN loss, or dedicated neural networks. Our method is highly flexible for adding on to already trained neural networks as it directly uses the output from the DL-based method, and doesn't require knowledge of the network or retraining.

Moreover, we want to point out that our method has room for improvements using various k-space-based parallel imaging schemes that have been vastly explored in the last decades [21, 24, 25]. We hope that our work may help to rediscover traditional MR physics-based parallel imaging methods that have been replaced by the surge of deep learning techniques. Moreover, we want to emphasize that with the benefit of improved textures and high-frequency details, this method may provide a path for increased adoption of DL for accelerated MR applications in the clinics.

Our codes will be released after the double-blind review process.

2. Related works

2.1. DL MRI reconstruction

Here, we will shortly review DL-based MR reconstruction methods based on unrolled neural networks inspired by CS. An MR scanner images a patient's anatomy by acquiring measurements in the frequency domain using multiple receiver coils. Each coil acquires k-space samples modulated by the sensitivity of the coil. For i-th coil, the acquired signal can be denoted as:

$$\mathbf{k}_i = \mathcal{DF}(S_i x) + \epsilon, \tag{1}$$

where $\mathcal{D}$ is an binary mask operator that selects acquired k-space locations, $\mathcal{F}$ is the Fourier transform operator, S_i is a complex-valued sensitivity map of the i-th coil.

Classical CS enables reconstruction of under-sampled k-space measurements by employing a regularization function. CS methods solve the following optimization problem:

$$\hat{\mathbf{x}} = \text{argmin}_{\mathbf{x}} \frac{1}{2} \sum_i \left\| \mathcal{D}\mathcal{F}\left(S_i \mathbf{x}\right) - \tilde{\mathbf{k}}_i \right\|^2 + \lambda \Psi(\mathbf{x}) \quad (2)$$

$$= \text{argmin}_{\mathbf{x}} \frac{1}{2} \|A\mathbf{x} - \tilde{\mathbf{k}}\|^2 + \lambda \Psi(\mathbf{x}), \quad (3)$$

where Ψ is a predetermined regularization function, A is a linear operator that is composed of sensitivity map projection, Fourier transform, and discrete sampling. $\tilde{\mathbf{k}}$ is a vector of acquired k-space. The goal of the optimization is to reconstruct an image x that best matches the measured data $\tilde{\mathbf{k}}$. The regularization function Ψ, and the regularization parameter λ incorporate image priors to help constrain the problem.

Unrolled neural networks go beyond CS by extending the regularization function to be data-adaptive and highly nonlinear. The regularization function can be set to have a CNN structure, and the iterative reconstruction algorithm to solve the optimization problem can be unrolled to a deep neural network, where the free parameters and the function can be learned by training end-to-end with a fully sampled training dataset.

While various architectures have been proposed [1, 7], the neural network operates on the coil-combined intermediate image-domain, and intermediate k-space output is untouched. Moreover, the loss is mostly determined by only the image domain, and the k-space output is not well examined or constrained.

2.2. K-space domain parallel imaging: SPIRiT

In order to better embrace the underlying principle of our refinement method, we will revisit SPIRiT [21], a well-established k-space domain parallel-MRI reconstruction method. SPIRiT reconstructs the parallel MRI image by enforcing the self-consistency (or neighborhood consistency) of multi-coil k-space data. It is based on the fact that each k-space data point of a given coil can be formulated as a linear combination of the multi-coil signals of its neighboring k-space points [20].

SPIRiT first estimates the linear relation of the intra- and inter-coil k-space data from a small area of the fully sampled k-space center, which is also called ACS, and then applies this relationship to the rest of the k-space data.

In an operator form, this can be expressed as:

$$\mathbf{K} = \mathcal{G}\mathbf{K} \quad (4)$$

where $\mathbf{K}$ denotes k-space data for all coils, and $\mathcal{G}$ is an operator that convolves the k-space data with a series of calibration kernels that are estimated from the ACS. Then, the reconstruction can be performed by solving an optimization problem given by,

$$\begin{array}{ll} \text{minimize} & \|(\mathcal{G} - I)k\|^2 \\ \text{s.t.} & \|\mathcal{D}k - y\|^2 \leq \epsilon, \end{array} \quad (5)$$

where k is the reconstructed k-space, y the acquired data and $\mathcal{D}$ is an operator that selects only the acquired k-space locations. Eq. 5 can be reformulated in the unconstrained Lagrangian form,

$$\text{argmin}_{\boldsymbol{k}} \|\mathcal{D}k - y\|^2 + \lambda \|(\mathcal{G} - I)k\|^2. \quad (6)$$

One can also include additional penalty term $\mathcal{R}(\cdot)$ in Eq. 3 in the objective function that expresses the prior knowledge in the reconstruction by,

$$\text{argmin}_{\boldsymbol{k}} \|\mathcal{D}k - y\|^2 + \lambda_1 \|(\mathcal{G} - I)k\|^2 + \lambda_2 \mathcal{R}(k). \quad (7)$$

For example, in the l_1-SPIRiT model, l_1 wavelet operator is chosen as the additional penalty term. Here, the SPIRiT kernel($\mathcal{G} - I$) is often referred also as the annihilating filter in the frequency domain, where the reconstruction can be performed by exploiting the null-space constraint for the missing k-space recovery.

Similar methods can also be used for refining or correcting k-space errors, such as ghosting artifacts in EPI [26], phase errors in multishot MRI [27] and k-space outliers [28] which may be analogous to our work.

3. Proposed refinement method

The refinement requires three inputs: acquired undersampled k-space, SPIRiT kernel that is derived from the ACS line of the undersampled k-space, and the DL reconstructed k-space. The overall scheme of the proposed refinement method and inputs required for the refinement are demonstrated in Fig. 1A.

Before delving into implementation details, we want to clarify that the main objective of the refinement is to further constrain the DL estimated k-space to satisfy the neighborhood relationship of GRAPPA/SPIRiT. For example in Fig. 2, the DL estimated k-space tends to contain structural artifacts as well as line-like discrepancies between acquired and estimated k-space lines. Enforcing the neighborhood constraint can reduce the artifacts and improve high-frequency k-space, as demonstrated by the refined k-space in Fig. 2.

First, we will explain how to obtain the SPIRiT kernel (denoted as auto-calibration in Fig. 1A). Prior to calibration, the virtual conjugate coil (VCC) concept [29] is employed to further double the number of channels. VCC is

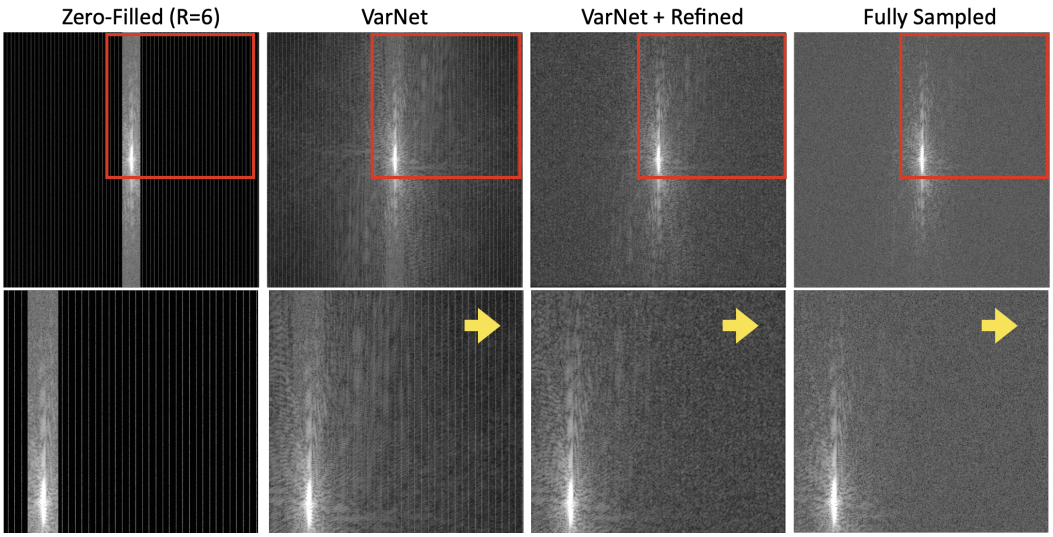

Figure 2. The top row (left to right) shows undersampled k-space with R=6, estimated k-space from Variational Network, refined k-space with the proposed method, and fully sampled k-space. The bottom row shows the magnified view of each result.

a technique to improve the conditioning of the SPIRiT reconstruction problem by exploiting the neighborhood constraint of the conjugate symmetric counterpart. The calibration region (orange box in Fig. 1A) is then converted to calibration matrix by sweeping the kernel, constructing block-Hankel structure [23]. The calibration matrix is then further denoised via singular value thresholding with a value suggested by [30]. The SPIRiT kernel can then be estimated via Tikhonov regularized least-squares, which has an analytic solution [21].

Second, we consider the following optimization problem:

$$\text{argmin}_{\boldsymbol{k}} \|\mathcal{D}k - y\|^2 + \lambda_1 \|(\mathcal{G} - I)k\|^2 + \lambda_2 \|\mathcal{D}_c(k - \tilde{k})\|^2, \tag{8}$$

where $\mathcal{D}_c$ is an operator that selects un-acquired k-space locations, $\tilde{k}$ is the DL-estimated k-space, assuming that the DL estimation has estimated k-space with some degree of precision in mean squared error.

The λ_1 and λ_2 parameters are tunable and can determine the level of image quality of the refinement. Increasing λ_1 puts more constraint on the k-space neighborhood relationship, and λ_2 emphasizes DL estimation. Typically, the accuracy of DL reconstruction decreases with increased acceleration factors, therefore we decreased λ_2 for high acceleration cases. In our experiments, we empirically selected $\lambda_1 = 5$, $\lambda_2 = 0.05$ for $R = 4$; and $\lambda_1 = 15$, $\lambda_2 = 0.01$ for $R = 6$.

For solving the optimization problem, we used the conjugate gradient (CG) algorithm with 300 iterations to ensure convergence. Pseudo-code for implementation is presented in Algorithm 1.

In current refinement, we have separated the two stages — reconstruction of the under-sampled k-space and the correction of the high-frequency k-space. Following optimization may also be considered if DL reconstruction network has been trained for general MR reconstruction for various accelerations such as plug-in-prior approach [31].

$$\text{argmin}_{\boldsymbol{k}} \|\mathcal{D}k - y\|^2 + \lambda_1 \|(\mathcal{G} - I)k\|^2 + \lambda_2 \|\mathcal{D}_c(k - DL(k))\|^2. \tag{9}$$

Algorithm 1 k-space refinement using SPIRiT

Inputs:
y = Acquired k-space,
$\tilde{k}$ = DL-estimated k-space
Output:
Refined k-space
1: **VCC:**
Apply virtual conjugate coil augmentation on $y, \tilde{k}$
2: **Calibrate:**
Calibrate SPIRIT kernel using y
$G - I \leftarrow$ Auto-calibration line
3: **Refine:**
minimize Eq.8 with CG

4. Experiments

In our experiments, we tested our method with two different datasets in order to explore the potential of our refinement method. We tested our refinement method on the Variational Network [32] available at https://github.com/facebookresearch/fastMRI. We will describe precise details in the following sections.

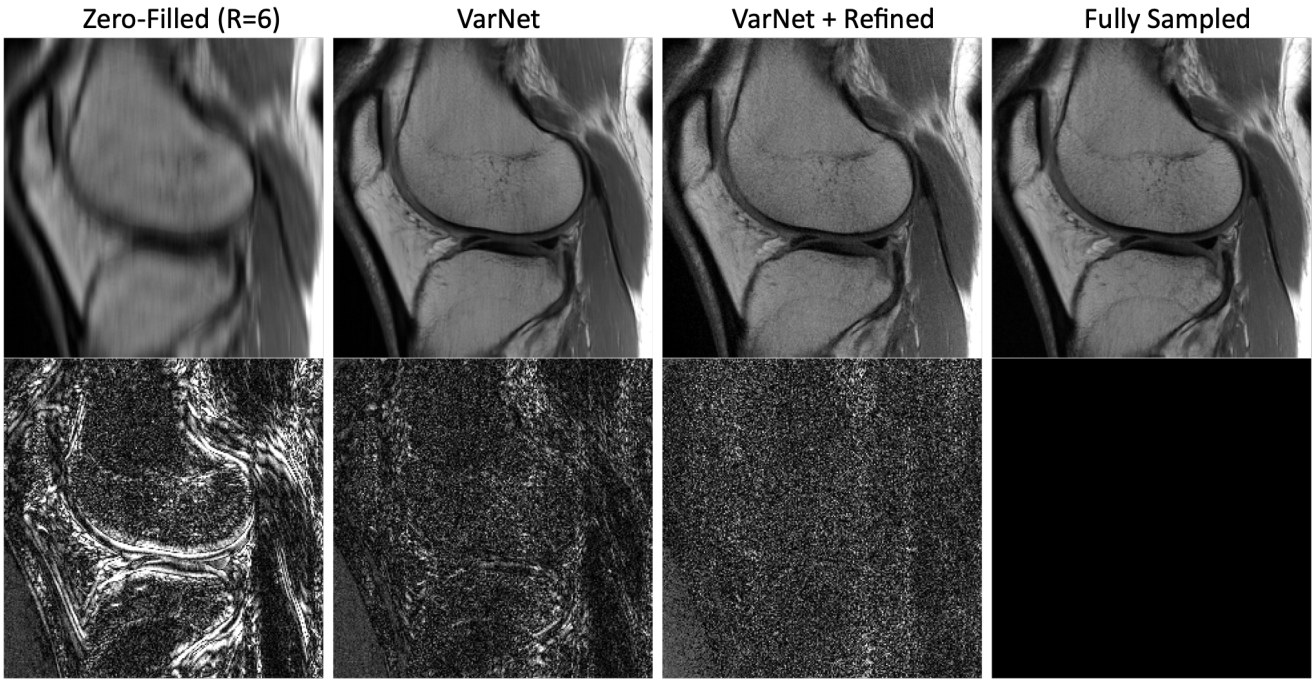

Figure 3. The top row (left to right) shows zero-filled reconstruction, Variational Network reconstruction, reconstruction from refinement, and reconstruction from fully-sampled k-space. The bottom row shows the error maps ($10\times$) corresponding to each reconstruction result above.

4.1. Dataset

In our experiments, we used two knee datasets described in [7], which are also shared at mridata.org. Two datasets include fully-sampled k-space of Sagittal Proton Density (PD) Weighted dataset (knee dataset-1), and Sagittal Fat-Saturated T2 Weighted dataset (knee dataset-2). In these datasets, each subject was scanned with a 15- channel knee coil. Sequence parameters are described in the paper [7].

Each dataset includes a total of 600 images from 20 subjects. We divided the datasets for training: 14 subjects, validation: 2 subjects, and testing: 4 subjects.

The sensitivity maps were computed from a block of size 24×24 using nonlinear inversion [33], which is available in SigPy library [34].

4.2. DL reconstruction

We used the k-space-based Variational Network (VarNet) as the baseline DL Reconstruction [32]. This network differs from end-to-end Variational Networks in that sensitivity maps are pre-calibrated and fed to the network. The reason for using this network instead of the original Variational Network [7] is that we want to leverage the improvements reported in the newer version, which was updating the soft data consistency directly in the k-space. U-Net was chosen as the CNN architecture in VarNet, and 10 unrolls were used.

We separately trained networks for the two datasets with equispaced sampling patterns, which samples l-th low frequency lines from the center of the k-space and every r-th line from the remaining k-space. We chose l, r randomly to match the overall acceleration. For each dataset, two networks corresponding to two acceleration factors $(R = 4, 6)$ were individually trained.

4.3. Refinement

As demonstrated in Algorithm 1, the inputs to the refinement are DL-reconstructed k-space and original undersampled k-space. For each slice, DL-reconstructed k-space was obtained by forward-propagating the undersampled k-space with the trained Variational network, and Fourier transforming the multi-coil reconstructed images right before the final coil-combination stage at the Variational Network.

For evaluation, we transformed the three multi-coil k-spaces (DL-reconstructed k-space, refined k-space, fully sampled k-space) to corresponding final images with inverse Fourier Transform and performed coil combination using root-sum-squares for each pixel:

$$\hat{\mathbf{x}} = RSS\left(\mathbf{x}_1, \ldots, \mathbf{x}_N\right) = \sqrt{\sum_{i=1}^{N} \left|\mathbf{x}_i\right|^2} \qquad (10)$$

			General IQM		High Freq IQM	
Contrast	Acceleration	Method	PSNR (dB) ↑	SSIM ↑	HFEN ↓	GMSD ↓
Sagittal PDw	4	VarNet	42.23	0.9767	0.3776	0.009974
		VarNet + Refine	42.29	0.974	0.3461	0.008474
	6	VarNet	37.97	0.9558	0.707	0.02791
		VarNet + Refine	37.75	0.9449	0.6278	0.02181
Sagittal T2w	4	VarNet	40.34	0.9687	0.4987	0.01757
		VarNet + Refine	40.33	0.9639	0.4188	0.0124
	6	VarNet	36.64	0.9385	0.9088	0.03948
		VarNet + Refine	37.22	0.9411	0.7656	0.03071

Table 1. Quantitative image quality assessment on the two test datasets for $R = 4, R = 6$. Average values of PSNR, SSIM, HFEN, and GMSD are reported. The higher the PSNR and SSIM, the better the image quality. The lower the HFEN and GMSD, the better.

4.4. Evaluation Metrics

To evaluate the impact of the proposed refinement on the high-frequency image details, we adopt two metrics: high-frequency error norm (HFEN) and Gradient Magnitude Standard Deviation (GMSD) [35]. The HFEN metric was used in DLMRI work [36] and GSMD in [37]. Both metrics were used to assess the reconstruction quality of edges and high-frequency image details.

The HFEN metric is formulated as the L2 norm of the difference between the filtered reconstructed image and filtered fully sampled image. It uses the Laplacian of Gaussian (LoG) filter. The GMSD metric is the standard deviation of the gradient magnitude similarity maps of the two images. The default parameters in [35, 36] were used for calculating the metrics. Widely used general image quality metrics (IQM)—i.e., PSNR (Peak Signal to Noise Ratio) and SSIM (Structural Similarity Index)—were also computed to assess the overall quality of the reconstructed images.

5. Results

The performance of the proposed refinement method was tested on the examples in the test dataset. Figure 2 compares the k-spaces of zero-filled reconstruction with a net acceleration of 6, Variational Network reconstruction, our proposed refinement, and the fully sampled k-space. For better visualization, the k-spaces only for the 1^{st} coil are displayed in log scale. We observed that the Variational Network reconstructions tend to produce under-estimation in the peripheral region (lower pixel values compared to fully sampled k-space), corresponding to high frequency band (Yellow arrow). Also, the DL estimated k-space contains line-like discrepancies between acquired and non-acquired k-space lines. As can be seen, the proposed refinement successfully recovers the underestimation and reduces the structural k-space error to some extent compared to DL-estimated k-space.

Figure 3 displays the effect of the refinement method in the final reconstructed image with a net acceleration of 6. We observe that the refinement method produces sharper images and better textures than the Variational Network. In the error map, the proposed method has errors that are more incoherent than the Variational Network results.

Measures of the four IQMs—i.e., PSNR, SSIM, HFEN, and GMSD—are summarized in Table 1 for the test dataset. For general IQMs such as PSNR and SSIM, the refined method shows competitive results compared to Variational Network. For high-frequency IQM (HFEN and GMSD) our method shows significantly improved values. Specifically, for all cases, HFEN decreased by nearly 10% and GSMD decreased by nearly 20%. This indicates that our method improves the high-frequency details.

Figure 4 displays the result of a slice containing ligament of the knee at the sagittal view of Proton Density weighted image (red box). As can be seen, the refinement method recovers the intricate texture of the ligament which is missing in the Variational Network result. Also, the refinement result shows sharper images with better texture details.

Figure 5 displays the result for a Fat-saturated T2-weighted image with a net acceleration of 4. It can be seen that occluded vessels—e.g., the yellow arrow on the lower left—and fine details of the knee joints—e.g., articular cartilage indicated by the yellow arrow on the upper right—have been recovered with the proposed method. Overall, the refined reconstruction produces images that are sharper and contain better textural details.

Figure 6 demonstrates multi-slice results within a single subject in the sagittal T2 weighted dataset for a net acceleration of 4. It can be seen that the proposed refinement method consistently improves high-frequency image details for different slice positions.

6. Discussions and Conclusion

This paper proposes a SPIRiT based refinement scheme that can be added to DL-based methods to improve high-frequency image quality and texture. Our results demon-

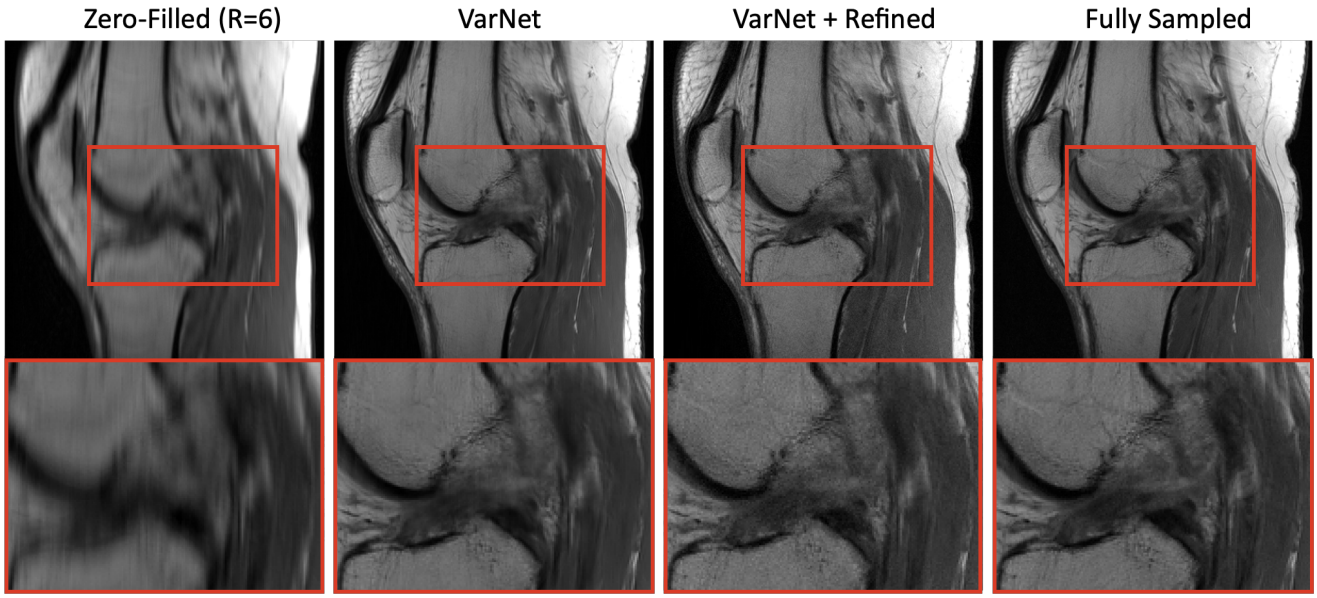

Figure 4. On the top row (left to right), zero-filled reconstruction, Variational Network reconstruction, reconstruction from refinement, and reconstruction from fully-sampled k-space are shown for an example slice from knee-dataset 1. The bottom row shows the magnified images for the ROI drawn in the images above.

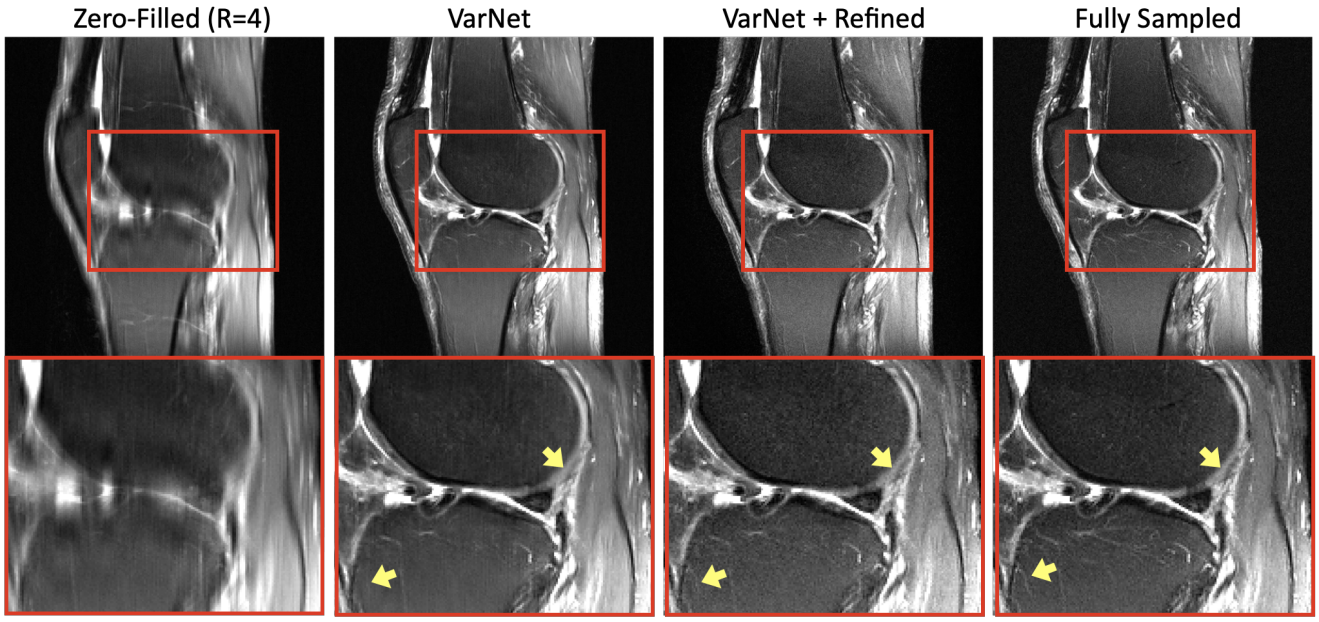

Figure 5. On the top row (left to right), zero-filled reconstruction, Variational Network reconstruction, reconstruction from refinement, and reconstruction from fully-sampled k-space are shown for an example slice from knee-dataset 2. The bottom row shows the magnified images for the ROI drawn in the images above.

strate that refining DL output with traditional k-space interpolation methods can improve DL reconstruction, and suggest that it may be valuable to consider using the proposed refinement with DL-based reconstructions.

The proposed method enhances high-frequency details of the Variational Network output both visually and quantitatively for the two datasets. It produces sharper images and better preserves subtle details such as microvascular structures and boundaries between different tissues. Some textures that are blurred in Variational Network reconstruc-

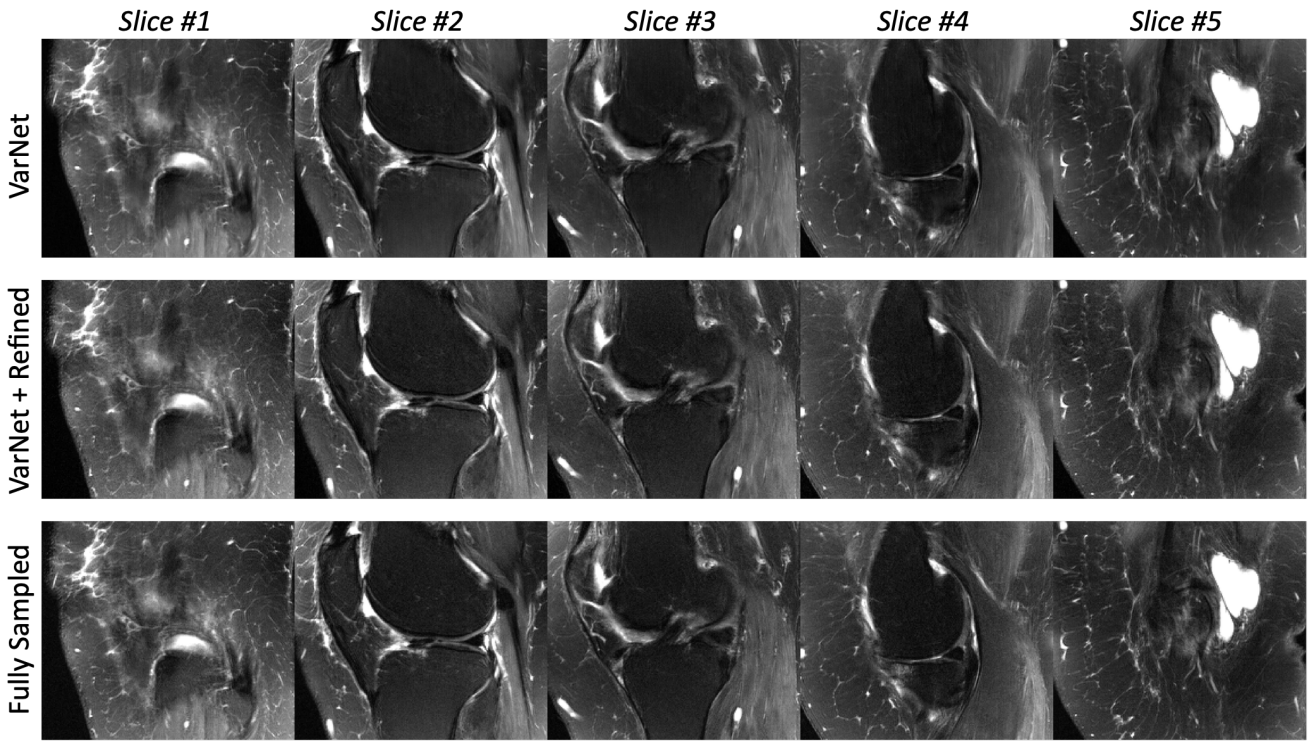

Figure 6. Results of slices from a single subject in the test set are shown. The first, second, and third rows show results of five slices from the Variational network, our refined method, and fully-sampled reference, respectively.

tion are recovered after refinement. IQMs indicate that the proposed method is highly effective for improving high-frequency image details without sacrificing general image quality metrics.

The constraint we used in this work was k-space neighborhood consistency, which can be easily obtained from the auto-calibration lines that are typically acquired in an MRI scan. Our refinement method uses the unique local neighborhood relationship of k-space [20, 21]. Exploring other constraints (e.g., finite image support, slowly varying phase) [25] may improve the refinement for certain MRI applications. Moreover, embedding such relationships in the network architecture or losses might regularize the networks to learn better reconstructions. These might be new research directions that future studies can focus on.

This method may have advantages in clinical settings such as in the identification of detailed ligaments and microscopic structures inside the joint due to the clearer tissue boundaries. Future study will include cases that has lesions such as meniscus tear and run blind tests between the existing and proposed methods assessed by multiple musculoskeletal radiologists.

One limitation of our method is the extended inference time. The current implementation is based on the naive SPIRiT implementation and requires about $20s$ for each slice. Using preconditioned conjugate descent algorithms and converting the convolution operator to image-domain multiplication will greatly accelerate our method. Moreover, as Eq.8 is in the form of a convex optimization problem, it is relatively much faster and more stable than non-convex optimization problems. Moreover, using advanced optimization schemes such as parallelization may be beneficial for further accelerating our method within a clinically feasible runtime.[38]. Moreover, thorough comparison studies with various UNN architectures, GAN-reconstruction, and dual-domain approaches will be needed to be performed to validate the efficacy of our method.

References

[1] Y. Yang, J. Sun, H. Li, and Z. Xu, "ADMM-Net: A Deep Learning Approach for Compressive Sensing MRI," *arXiv*, 2017.

[2] 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, pp. 487–492, 2018.

[3] T. Eo, Y. Jun, T. Kim, J. Jang, H. 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.

[4] M. Mardani, E. Gong, J. Y. Cheng, S. S. Vasanawala, G. Zaharchuk, L. Xing, and J. M. Pauly, "Deep Generative Adversarial Neural Networks for Compressive Sensing MRI," *IEEE Transactions on Medical Imaging*, vol. 38, no. 1, pp. 167–179, 2019.

[5] G. Yang, S. Yu, H. Dong, G. Slabaugh, P. L. Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, Y. Guo, and D. Firmin, "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, 2017.

[6] Y. Han, L. Sunwoo, and J. C. Ye, "${k}$ -Space Deep Learning for Accelerated MRI," *IEEE Transactions on Medical Imaging*, vol. 39, no. 2, pp. 377–386, 2020.

[7] K. Hammernik, T. Klatzer, E. Kobler, M. P. Recht, D. K. Sodickson, T. Pock, and F. Knoll, "Learning a variational network for reconstruction of accelerated MRI data," *Magnetic Resonance in Medicine*, vol. 79, no. 6, pp. 3055–3071, 2018.

[8] 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.

[9] F. Knoll, T. Murrell, A. Sriram, N. Yakubova, J. Zbontar, M. Rabbat, A. Defazio, M. J. Muckley, D. K. Sodickson, C. L. Zitnick, and M. P. Recht, "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.

[10] M. J. Muckley, B. Riemenschneider, A. Radmanesh, S. Kim, G. Jeong, J. Ko, Y. Jun, H. Shin, D. Hwang, M. Mostapha, S. Arberet, D. Nickel, Z. Ramzi, P. Ciuciu, J.-L. Starck, J. Teuwen, D. Karkalousos, C. Zhang, A. Sriram, Z. Huang, N. Yakubova, Y. W. Lui, and F. Knoll, "Results of the 2020 fastMRI Challenge for Machine Learning MR Image Reconstruction," *IEEE Transactions on Medical Imaging*, vol. PP, no. 99, pp. 1–1, 2021.

[11] C. M. Sandino, J. Y. Cheng, F. Chen, M. Mardani, J. M. Pauly, and S. S. Vasanawala, "Compressed Sensing: From Research to Clinical Practice With Deep Neural Networks," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 117–127, 2020.

[12] F. Knoll, K. Hammernik, C. Zhang, S. Moeller, T. Pock, D. K. Sodickson, and M. Akakaya, "Deep-Learning Methods for Parallel Magnetic Resonance Image Reconstruction," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 128–140, 2020.

[13] D. Liang, J. Cheng, Z. Ke, and L. Ying, "Deep Magnetic Resonance Image Reconstruction," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 141–151, 2020.

[14] K. Hammernik, J. Schlemper, C. Qin, J. Duan, R. M. Summers, and D. Rueckert, "Systematic evaluation of iterative deep neural networks for fast parallel MRI reconstruction with sensitivity-weighted coil combination," *Magnetic Resonance in Medicine*, 2021.

[15] J. P. Cohen, M. Luck, and S. Honari, "Distribution Matching Losses Can Hallucinate Features in Medical Image Translation," *arXiv*, 2018.

[16] V. Edupuganti, M. Mardani, S. Vasanawala, and J. Pauly, "Uncertainty Quantification in Deep MRI Reconstruction," *IEEE Transactions on Medical Imaging*, vol. 40, no. 1, pp. 239–250, 2021.

[17] K. Wang, J. I. Tamir, S. X. Yu, and M. Lustig, "High-Fidelity Reconstruction with Instance-wise Discriminative Feature Matching Loss," in *Proceedings of the 27th Annual Meeting of ISMRM*, 2021.

[18] N. M. Singh, J. E. Iglesias, E. Adalsteinsson, A. V. Dalca, and P. Golland, "Joint Frequency and Image Space Learning for Fourier Imaging," *arXiv*, 2020.

[19] B. Zhou and S. K. Zhou, "DuDoRNet: Learning a Dual-Domain Recurrent Network for Fast MRI Reconstruction with Deep T1 Prior," *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 00, pp. 4272–4281, 2020.

[20] 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.

[21] M. Lustig and J. M. Pauly, "SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space," *Magnetic Resonance in Medicine*, vol. 64, no. 2, pp. 457–471, 2010.

[22] J. P. Halder, "Autocalibrated Loraks for Fast Constrained MRI Reconstruction," *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pp. 910–913, 2015.

[23] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and M. Lustig, "ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA," *Magnetic Resonance in Medicine*, vol. 71, no. 3, pp. 990–1001, 2014.

[24] K. H. Jin, D. Lee, and J. C. Ye, "A Novel K-Space Annihilating Filter Method for Unification Between Compressed Sensing and Parallel MRI," *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pp. 327–330, 2015.

[25] J. P. Haldar and J. Zhuo, "P-LORAKS: Low-rank modeling of local k-space neighborhoods with parallel imaging data," *Magnetic Resonance in Medicine*, vol. 75, no. 4, pp. 1499–1514, 2016.

[26] R. A. Lobos, W. S. Hoge, A. Javed, C. Liao, K. Setsompop, K. S. Nayak, and J. P. Haldar, "Robust autocalibrated structured low-rank EPI ghost correction," *Magnetic Resonance in Medicine*, vol. 85, no. 6, pp. 3403–3419, 2021.

[27] M. Mani, M. Jacob, D. Kelley, and V. Magnotta, "Multi-shot sensitivity-encoded diffusion data recovery using structured low-rank matrix completion (MUSSELS)," *Magnetic Resonance in Medicine*, vol. 78, no. 2, pp. 494–507, 2017.

[28] K. H. Jin, J. Um, D. Lee, J. Lee, S. Park, and J. C. Ye, "MRI artifact correction using sparse + low-rank decomposition of annihilating filter-based hankel matrix," *Magnetic Resonance in Medicine*, vol. 78, no. 1, pp. 327–340, 2017.

[29] M. Blaimer, M. Gutberlet, P. Kellman, F. A. Breuer, H. Köstler, and M. A. Griswold, "Virtual coil concept for improved parallel MRI employing conjugate symmetric signals," *Magnetic Resonance in Medicine*, vol. 61, no. 1, pp. 93–102, 2009.

[30] M. Gavish and D. L. Donoho, "The Optimal Hard Threshold for Singular Values is $4/\sqrt{3}$," *IEEE Transactions on Information Theory*, vol. 60, no. 8, pp. 5040–5053, 2014.

[31] A. P. Yazdanpanah, O. Afacan, and S. Warfield, "Deep plug-and-play prior for parallel mri reconstruction," *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, pp. 3952–3958, 2019.

[32] A. Sriram, J. Zbontar, T. Murrell, A. Defazio, C. L. Zitnick, N. Yakubova, F. Knoll, and P. Johnson, "End-to-End Variational Networks for Accelerated MRI Reconstruction," *arXiv*, 2020.

[33] M. Uecker, T. Hohage, K. T. Block, and J. Frahm, "Image reconstruction by regularized nonlinear inversion—Joint estimation of coil sensitivities and image content," *Magnetic Resonance in Medicine*, vol. 60, no. 3, pp. 674–682, 2008.

[34] F. Ong and M. Lustig, "Sigpy: a python package for high performance iterative reconstruction," *Proceedings of the International Society of Magnetic Resonance in Medicine, Montréal, QC*, vol. 4819, 2019.

[35] W. Xue, L. Zhang, X. Mou, and A. C. Bovik, "Gradient magnitude similarity deviation: A highly efficient perceptual image quality index," *IEEE Transactions on Image Processing*, vol. 23, no. 2, pp. 684–695, 2013.

[36] S. Ravishankar and Y. Bresler, "MR Image Reconstruction from Highly Undersampled k-Space Data by Dictionary Learning," *IEEE Transactions on Medical Imaging*, vol. 30, no. 5, pp. 1028–1041, 2011.

[37] A. Mason, J. Rioux, S. E. Clarke, A. Costa, M. Schmidt, V. Keough, T. Huynh, and S. Beyea, "Comparison of Objective Image Quality Metrics to Expert Radiologists' Scoring of Diagnostic Quality of MR Images," *IEEE Transactions on Medical Imaging*, vol. 39, no. 4, pp. 1064–1072, 2020.

[38] M. Murphy, M. Alley, J. Demmel, K. Keutzer, S. Vasanawala, and M. Lustig, "Fast l1-spirit compressed sensing parallel imaging mri: scalable parallel implementation and clinically feasible runtime," *IEEE transactions on medical imaging*, vol. 31, no. 6, pp. 1250–1262, 2012.