Published in final edited form as:

IEEE J Sel Top Signal Process. 2020 October ; 14(6): 1280–1291. doi:10.1109/jstsp.2020.3003170.

Dense Recurrent Neural Networks for Accelerated MRI: History-Cognizant Unrolling of Optimization Algorithms

Seyed Amir Hossein Hosseini [Student Member, IEEE],

Department of Electrical and Computer Engineering, and Center for Magnetic Resonance Research, University of Minnesota, Minneapolis, MN, 55455

Burhaneddin Yaman [Student Member, IEEE],

Department of Electrical and Computer Engineering, and Center for Magnetic Resonance Research, University of Minnesota, Minneapolis, MN, 55455

Steen Moeller,

Center for Magnetic Resonance Research, University of Minnesota, Minneapolis, MN, 55455

Mingyi Hong [Member, IEEE],

Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN, 55455

Mehmet Akçakaya [Member, IEEE]

Department of Electrical and Computer Engineering, and Center for Magnetic Resonance Research, University of Minnesota, Minneapolis, MN, 55455

Abstract

Inverse problems for accelerated MRI typically incorporate domain-specific knowledge about the forward encoding operator in a regularized reconstruction framework. Recently physics-driven deep learning (DL) methods have been proposed to use neural networks for data-driven regularization. These methods unroll iterative optimization algorithms to solve the inverse problem objective function, by alternating between domain-specific data consistency and data-driven regularization via neural networks. The whole unrolled network is then trained end-to-end to learn the parameters of the network. Due to simplicity of data consistency updates with gradient descent steps, proximal gradient descent (PGD) is a common approach to unroll physics-driven DL reconstruction methods. However, PGD methods have slow convergence rates, necessitating a higher number of unrolled iterations, leading to memory issues in training and slower reconstruction times in testing. Inspired by efficient variants of PGD methods that use a history of the previous iterates, we propose a history-cognizant unrolling of the optimization algorithm with dense connections across iterations for improved performance. In our approach, the gradient descent steps are calculated at a trainable combination of the outputs of all the previous regularization units. We also apply this idea to unrolling variable splitting methods with quadratic relaxation. Our results in reconstruction of the fastMRI knee dataset show that the proposed history-cognizant approach reduces residual aliasing artifacts compared to its conventional

unrolled counterpart without requiring extra computational power or increasing reconstruction time.

Keywords

Inverse problems; unrolled optimization algorithms; physics-driven deep learning; neural networks; MRI reconstruction; recurrent neural networks

I. Introduction

Inverse problems have been extensively utilized in many imaging modalities, including magnetic resonance imaging (MRI) [1–7]. These inverse problems are directly driven from the physics of data acquisition, known as the forward operator, which incorporates domain-specific knowledge. Such inverse problems are typically ill-conditioned and thus regularizers are incorporated into the respective objective functions. Subsequently, iterative algorithms are employed to solve the regularized inverse problem by alternating between enforcing physics-driven data consistency and regularization using preselected priors.

Recently, deep learning (DL) has emerged as an alternative state-of-the-art technique for solving such inverse problems in imaging applications [8–39]. Several of the earlier works concentrated on training convolutional neural networks (CNN) to solve the problem without incorporating explicit knowledge of the acquisition physics, typically requiring re-training when acquisition parameters are modified [8–18]. Another line of work focuses on a physics-driven formulation, where the known forward model is incorporated into training, utilizing domain knowledge to solve the corresponding inverse problem [19–35]. In physics-driven approaches, a conventional iterative optimization algorithm for solving a regularized least-squares problem is unrolled for a fixed number of iterations. This unrolled network alternates between data consistency and regularization steps similar to the conventional algorithm, where the regularizers are implemented via neural networks. The unrolled network is trained end-to-end to learn the network parameters, which predominantly characterize the regularization units.

There are numerous approaches for solving regularized least squares problems, which have been explored for algorithm unrolling in physics-driven DL reconstruction methods [19–35]. These methods include gradient descent (GD) [19], proximal gradient descent (PGD) [20–23], variable-splitting (VS) [24–31] and primal-dual (PD) [32, 33] methods. The GD, PGD and PD unrolling typically use a gradient descent step for data consistency, while for VS-based methods, either a gradient descent [26] or conjugate gradient (CG) data consistency can be used, with the latter being utilized when the least squares problem involving the forward encoding operator becomes computationally expensive [24, 25, 27–31]. Data consistency updates with gradient descent step is generally computationally inexpensive. However, these optimization methods have slower convergence, which necessitates higher number of unrolled iterations for improved performance, which in turn can lead to numerical issues such as gradient vanishing and memory constraints.

In this study, we seek to develop a new methodology for unrolling inverse problem optimization algorithms into dense recurrent neural networks (RNN) for improved reconstruction performance. Motivated by accelerated PGD approaches, which calculate the gradient descent at a combined history of previous iterations rather than the immediate previous iteration only [40–47], we propose a history-cognizant unrolling of the optimization algorithms with skip connections across unrolled iterations, leading to a dense RNN architecture. Skip connections have been previously shown to facilitate information flow through neural networks in architectures such as ResNet [48] and DenseNet [49]. Different than these previous neural network design problems, we propose to utilize such connections in an unrolled RNN architecture for solving inverse problems, which readily comes with its domain-specific design considerations. We compare the conventional unrolling and the proposed history-cognizant unrolling of the PGD algorithm in accelerated multi-coil MRI reconstruction using the fastMRI knee dataset [50], showing both visual and quantitative improvement in reconstruction quality with the proposed methodology. We also extend the same unrolling strategy to other optimization algorithms and investigate the performance improvement.

II. Theory

In this paper, boldface lowercase symbols denote vectors, boldface uppercase symbols denote matrices and $(\cdot)^H$ represents the matrix conjugate transpose operation.

A. Forward Model and Inverse Problem

In a multi-coil MRI system with n_c coils, let $\mathbf{y}$ be the undersampled noisy data from all coils and $\mathbf{x}$ be the underlying image to be recovered. The forward linear operator in this system, $\mathbf{A}: \mathbb{C}^{M \times N} \rightarrow \mathbb{C}^P$, includes sensitivities of the receiver coil array and a partial Fourier matrix for undersampling in k-space, i.e. spatial Fourier domain [1]. Thus, the forward model is given as

$$\mathbf{y} = \mathbf{Ax} + \mathbf{n}, \quad (1)$$

where $\mathbf{n}$ is measurement noise. The corresponding inverse problem is typically ill-conditioned. Hence, regularization is frequently used, leading to a regularized least-squares method:

$$\arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{Ax}\|_2^2 + \mathcal{R}(\mathbf{x}), \quad (2)$$

where the first term enforces consistency with measurement data and $\mathcal{R}(\cdot)$ denotes a regularizer. In classical reconstruction approaches, regularizers have included Tikhonov [1, 2], edge-preserving functions [51] or sparsity-promoting terms in predefined domains [3–6]. A variety of algorithms can iteratively solve the optimization problem in (2) such as PGD [52], VS [53–55] and PD [56] methods.

B. Physics-Based DL Reconstruction by Unrolling PGD

Several physics-based DL methods have unrolled a PGD-based method [20–23] for a fixed number of iterations. Such PGD-based unrolling alternates between a proximal operation that is implicitly defined by a learned neural network, and a gradient descent step to enforce data consistency, as follows:

$$\begin{aligned} \mathbf{z}^{(i)} &= \text{Prox}_{\mathcal{R}}(\mathbf{x}^{(i-1)}) = \\ &= \arg \min_{\mathbf{z}} \|\mathbf{x}^{(i-1)} - \mathbf{z}\|_2^2 + \mathcal{R}(\mathbf{z}) \end{aligned} \quad (3a)$$

$$\mathbf{x}^{(i)} = \mathbf{z}^{(i)} + \mu_i \mathbf{A}^H (\mathbf{y} - \mathbf{A} \mathbf{z}^{(i)}) \quad (3b)$$

where μ_i is the gradient descent step size at unrolled iteration i , and the proximal operation in (3a) is performed by a neural network. This leads to the schematic in Figure 1a, where the unrolled RNN consists of multiple blocks of proximal (P) and data consistency (DC) units. The former is a neural network and the latter implements a gradient descent step based on the known forward encoding operator. The unrolled RNN is then trained end-to-end to learn the mapping from the acquired sensor-domain data to reference images for medical imaging reconstruction.

C. Proposed History-Cognizant PGD Unrolling with Dense Recurrent Neural Networks

One of the major appeals of the PGD unrolling scheme is its computational simplicity. However, PGD algorithm itself has a slow convergence rate [51], potentially requiring a large number of iterations (i.e., a large number of blocks in Figure 1a) when unrolled into a neural network for sufficient reconstruction quality. In turn, this increases memory requirements and impedes training complexity. On the other hand, several methods have been proposed to accelerate convergence rate of PGD [40–47]. The underlying theme for these methods is to calculate the gradient step not only at the most recent proximal operation output, but at a linear combination of all past such outputs.

We propose to adapt this general accelerated PGD approach for network unrolling. In this setting, the accelerated PGD methods involve the following steps:

$$\mathbf{z}^{(i)} = \text{Prox}_{\mathcal{R}}(\mathbf{x}^{(i-1)}) \quad (4a)$$

$$\mathbf{v}^{(i)} = \mathbf{F}(\mathbf{z}^{(i)}, \dots, \mathbf{z}^{(1)}) \quad (4b)$$

$$\mathbf{x}^{(i)} = \mathbf{v}^{(i)} + \mu_i \mathbf{A}^H (\mathbf{y} - \mathbf{A} \mathbf{v}^{(i)}), \quad (4c)$$

where $\mathbf{F}(\cdot)$ forms a linear combination of the previous proximal operator outputs. As a special case, Nesterov's method uses the previous two outputs [41], while more general versions have also been studied, in which all the past iterations are used [42, 46, 47].

For our unrolled network, we propose to learn $\mathbf{F}(\cdot)$ during training instead of specifying a pre-determined form for the linear combination of previous iterates $\mathbf{F}(\cdot)$ as in Nesterov's method [41]. This history-cognizant unrolling of the accelerated PGD algorithms leads to a dense recurrent neural network architecture (Dense-RNN) with skip connections across all unrolled iterations, as depicted in Figure 1b. In addition to the P and DC units to perform proximal operation (or regularization) and data consistency updates, each unrolled iteration also consists of a 1×1 convolutional layer implementing $\mathbf{F}(\cdot)$ on the concatenation (denoted by C) of all past proximal operator outputs prior to the DC unit. We refer to this method as history-cognizant PGD (HC-PGD) in the rest of this study.

III. Multi-Coil MRI Reconstruction Experiments and Results

The proposed history-cognizant algorithm unrolling approach was compared to the conventional algorithm unrolling in multi-coil MRI reconstruction. The input $\mathbf{x}_0$ to the unrolled network in this system is the zero-filled SENSE-1 [57] image, which is given as $\mathbf{A}^H \mathbf{y}$ for $\mathbf{A}$ including normalized coil sensitivities.

A. Knee MRI Datasets

Knee MRI datasets were obtained from the New York University (NYU) fastMRI initiative database [50] to investigate accelerated MRI reconstruction using conventional unrolled RNN and Dense-RNN architectures. Data were acquired without any acceleration on a clinical 3T system (Magnetom Skyra, Siemens, Germany) with a 15-channel knee coil. Data acquisition was performed using 2D turbo spin-echo sequences in coronal orientation with proton-density (Coronal PD) and proton-density weighted with fat suppression (Coronal PD-FS) weightings. Relevant imaging parameters were: resolution = $0.49 \times 0.44 \text{ mm}^2$, slice thickness = 3mm, matrix size = 320×368 for both datasets.

All data were uniformly sub-sampled at an acceleration rate of 4 by keeping the central 24 lines using the mask provided in the fastMRI database [50] retrospectively. Uniform undersampling for 2D acquisitions is typically considered a more challenging setup for DL reconstruction than random undersampling [19]. A 24×24 central window was used to generate coil sensitivity maps using ESPIRiT [58]. The training sets consisted of 300 slices from 15 subjects for each sequence. Testing was performed on 10 different subjects for each sequence.

B. Implementation Details

The optimization algorithms were unrolled for $T = 10$ iterations. The proposed HC-PGD algorithm was implemented using the Dense-RNN architecture in Figure 1b. Proximal operator unit outputs are first concatenated and then combined using convolutional layers with a 1×1 kernel size and 2 output channels (real and imaginary components) before being inputted to the subsequent data consistency unit. On the other hand, the PGD algorithm was implemented using a conventional architecture that only includes P and DC units, without skip connections, concatenation of previous iterates or 1×1 convolutions.

The same CNN was used to implement the proximal operation for both PGD and HC-PGD. A residual network (ResNet) that has been established in a different regression problem was

utilized [59]. This CNN consisted of 15 sub-residual blocks each having two convolutional layers of kernel size = 3×3 and output channels = 64. The first and second convolutional layers in each sub-residual block are followed by a ReLU activation function and a scaling factor of 0.1, respectively. In addition to sub-residual blocks, two input and output convolutional layers without any nonlinear activation function are to match number of desired channels. For both algorithm unrolling schemes, the proximal units shared parameters across unrolled iterations [27], leading to a total of 592,138 and 592,358 parameters for PGD and HC-PGD, respectively. In addition to network parameters, the gradient descent step sizes (μ_i) were also learned during training.

All networks were trained end-to-end, as summarized in Figure 2, by using Adam optimizer to minimize

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(\mathbf{x}_{ref}^i, f(\mathbf{y}^i, \mathbf{A}^i; \theta)), \quad (5)$$

where $\mathbf{x}_{ref}^i$, $\mathbf{y}^i$ and $\mathbf{A}^i$ are the reference SENSE-1 image, undersampled data and forward encoding operator for slice i , respectively and N is the number of training slices. $f(\mathbf{y}^i, \mathbf{A}^i; \theta)$ denotes the network output for slice i with the corresponding measurements and forward encoding operator, while θ contains the trainable parameters of the networks. A normalized ℓ_1 - ℓ_2 loss [39] was used for $\mathcal{L}(\cdot, \cdot)$:

$$\mathcal{L}(\mathbf{u}_1, \mathbf{u}_2) = \frac{\|\mathbf{u}_1 - \mathbf{u}_2\|_2}{\|\mathbf{u}_1\|_2} + \frac{\|\mathbf{u}_1 - \mathbf{u}_2\|_1}{\|\mathbf{u}_1\|_1}, \quad (6)$$

with $\mathbf{u}_1$ and $\mathbf{u}_2$ representing reference fully-sampled and network output images, respectively. Normalization of the ℓ_1 and ℓ_2 loss terms facilitates adjusting the ratio of contribution from each term which is equal after normalization in this case [39]. Training was performed over 100 epochs. All algorithms were implemented using TensorFlow v1.14 in Python v3.6, and processed on a workstation with an Intel E5-2640V3 CPU (2.6GHz and 256GB memory), and two NVIDIA Tesla V100 GPUs with 32GB memory.

C. Performance of HC-PGD versus PGD Unrolling

Experiments were conducted on both Coronal PD and PD-FS datasets to compare the reconstruction quality of the conventional PGD and the proposed HC-PGD unrolling with the Dense-RNN architecture. A total of 380 slices from 10 subjects (excluded from training) were utilized for testing. Reconstructed images were evaluated quantitatively using peak signal to noise ratio (PSNR) and structural similarity index (SSIM) metrics with respect to fully-sampled reference image.

Figure 3 depicts a representative slice from the Coronal PD dataset reconstructed using DL-MRI with PGD and HC-PGD unrolling methods. The first row shows the reference fully-sampled image, as well as the zero-filled acquired image, which serves as the input $\mathbf{x}^{(0)}$ to the network. Difference images with respect to the reference image, scaled by a factor of 10, are also provided to facilitate comparison by showing errors. DL-MRI based on the

proposed HC-PGD unrolling was able to suppress the strong aliasing artifacts successfully, while residual aliasing is visible on the DL-MRI reconstruction with conventional PGD unrolling. For all methods, there is a slight degree of blurring in final images, consistent with previous reports on DL-MRI methods with similar training loss functions [19]. Quantitative evaluations, shown in the lower left corner of the images, confirm the improvement using the proposed history-cognizant unrolling with Dense-RNN architecture compared to conventional algorithm unrolling.

The reconstructed images for a representative slice from the coronal PD-FS dataset, the difference images scaled by a factor of 10, as well as reference and zero-filled images are shown in Figure 4. There are slight residual aliasing artifacts in the DL-MRI reconstruction with PGD unrolling, which are removed further by the proposed HC-PGD unrolling. The improvement in quantitative metrics also align with these observations.

Figure 5 summarizes the quantitative analysis of the DL-MRI reconstructions using PSNR and SSIM metrics for the PGD and HC-PGD unrollings for both the coronal PD and PD-FS datasets. The box plots depict the interquartile range with the center at the median of the metric for the corresponding reconstruction method. Both metrics are improved over the whole test sets when using the proposed history-cognizant algorithm unrolling over the conventional unrolling approach. Further statistical tests were performed using the Wilcoxon signed rank test with a significance level of $P < 0.05$, indicating that all differences are statistically significant based on both metrics and in both datasets.

D. Insensitivity of Dense-RNN Improvement to the Regularization CNN Choice

In order to verify that the improvements between PGD and HC-PGD unrolling are not limited to the choice of CNN used for the proximal operation, the two approaches in Figure 1 were also implemented using a U-Net architecture for the regularization CNN. This was based on the U-net used in [50], which includes contracting and expanding paths up to 256 channels without batch normalization. Figure 6 depicts the results of this experiment on the coronal PD datasets, where the proposed history-cognizant unrolling outperforms the conventional unrolling methodology both visually and quantitatively. This highlights that the improvement from Dense-RNN is not limited to a specific choice of network architecture for the proximal operation.

E. Performance Evaluation for Random Undersampling

In order to evaluate performance of the proposed history-cognizant unrolling approach in random undersampling, we utilized an observation a previous study [19, 60], showing that physics-based networks trained on uniformly undersampled data generalizes to random undersampling patterns at the same rate, even though the opposite is not necessarily true. Thus, we employed the networks trained on uniformly undersampled data, as detailed in Section III-B, to reconstruct test slices which were randomly undersampled at 4-fold acceleration. The corresponding results, shown in Figure 7 indicate that the improvement trend remains similar to the uniform sampling scenario with the proposed HC-PGD method outperforming the conventional PGD in suppressing residual artifacts in both Coronal PD and PD-FS datasets, and improving the quantitative evaluation metrics.

F. Extensions to Other Optimization Algorithms

While PGD unrolling is commonly used for DL-MRI reconstruction [20–23], alternative algorithms have also been utilized [24–31]. One such approach relies on VS methods [24–31] for solving the objective function in Equation (2), which decouples data consistency and regularization terms into two blocks of variables [51] by introducing auxiliary variables to formulate a constrained objective function. This constrained problem is then relaxed to an unconstrained problem by enforcing similarity between these variables either via a quadratic penalty [26–31] or via an augmented Lagrangian approach [24, 25].

In the case of variable splitting with quadratic penalty (VSQP), the following objective function is used:

$$\arg \min_{\mathbf{x}, \mathbf{z}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \mathcal{R}(\mathbf{z}) + \beta \|\mathbf{x} - \mathbf{z}\|_2^2, \quad (7)$$

where $\mathbf{z}$ is an auxiliary variable and β is the quadratic penalty parameter. When this algorithm is unrolled as a neural network [26–31], reconstruction alternates between a proximal operation that is again learned by a neural network and data consistency, similar to PGD unrolling, as follows:

$$\mathbf{z}^{(i)} = \text{Prox}_{\mathcal{R}}(\mathbf{x}^{(i-1)}) \quad (8a)$$

$$\mathbf{x}^{(i)} = (\mathbf{A}^H \mathbf{A} + \mu \mathbf{I})^{-1} (\mathbf{A}^H \mathbf{y} + \beta \mathbf{z}^{(i)}) \quad (8b)$$

where regularization update in (8a) is performed by a neural network and β was absorbed into the regularizer for notational convenience. The data consistency update in (8b) is computationally challenging in many practical scenarios, such as multi-coil MRI reconstruction. Thus, this data consistency update can be solved iteratively via conjugate gradient (CG) to avoid matrix inversion [27].

Another line of work utilizes an augmented Lagrangian relaxation, leading to the unrolling of the ADMM algorithm [24, 25]. The augmented Lagrangian approach used for relaxation in ADMM crucially leads to an extra intermediate step to optimize the Lagrangian parameters in addition to the regularization and data consistency updates in VSQP unrolling. This intermediate step in general improves the convergence rate of VS methods, but also leads to more memory usage which can limit the number of unrolls that can be performed with limited GPU resources.

As an extension to the HC-PGD approach, both VSQP and ADMM methods may be unrolled in a history-cognizant manner using the proposed Dense-RNN in Figure 1, building on the optimization approaches in [44, 45]. These will be referred to as the history-cognizant VSQP (HC-VSQP) unrolling and the history-cognizant ADMM (HC-ADMM) unrolling respectively. We note that the main difference between HC-PGD and HC-VSQP unrolling is that the DC units in HC-VSQP are implemented with an unrolled CG algorithm for a different least squares problem instead of a gradient step. HC-ADMM further includes the Lagrangian parameter update as in ADMM-Net [24, 25].

To test the applicability of the history-cognizant unrolling scheme to different baseline algorithms, DL-MRI reconstruction based on VSQP, HC-VSQP, ADMM and HC-ADMM unrollings were also implemented on the coronal PD and PD-FS datasets, using the setup described in Section III-B. In each data consistency unit, CG was implemented with 10 internal iterations. Furthermore, due to the higher memory usage of the Lagrangian update in ADMM and HC-ADMM unrollings, these were unrolled for 8 iterations due to GPU memory limitations.

Figure 8 and Figure 9 depict representative slices from the Coronal PD and Coronal PD-FS datasets, respectively, reconstructed using DL-MRI with VSQP, ADMM, HC-VSQP and HC-ADMM unrolling methods, along with the error images scaled by a factor of 10, as well as the reference and zero-filled images. The reconstruction from the proposed history-cognizant unrolling approaches, i.e. HC-VSQP and HC-ADMM, successfully mitigate the residual artifacts apparent in the conventional VSQP and ADMM results, shown by the yellow arrows. Improvements are also observed in the corresponding quantitative metrics.

IV. Discussion

In this study, we proposed a history-cognizant method for unrolling optimization problems to solve regularized inverse problems in accelerated MRI using physics-based DL reconstruction. The proposed unrolling method builds on optimization approaches that utilize a history of the previous iterates via their linear combinations in updating the current estimate [40–45]. While these optimization procedures have been traditionally proposed to improve convergence rates or provide stronger theoretical guarantees, in the context of algorithm unrolling for physics-based DL reconstruction, they are used to enhance reconstruction quality without incurring additional computational complexity. Our method was predominantly motivated by accelerated PGD approaches [40–45] and was shown to improve upon conventional PGD unrolling for multi-coil MRI reconstruction for multiple scenarios. We also explored how the proposed unrolling can be applied to VS methods for solving a regularized least squares problem motivated by results from optimization theory [44, 45], and showed that it can improve over the conventional unrolling approach.

Our work makes several novel contributions to domain enriched learning for medical imaging, specifically from the perspective of signal processing. Our main contribution is the adoption of analytical insights for improving convergence from optimization theory to improve algorithm unrolling for deep learning reconstruction. Specifically, these insights were used to add skip connections in the unrolled neural network in a principled manner, termed history-cognizant unrolling. In turn, this led to a dense RNN architecture in a regression problem for the first time, while maintaining interpretability due to the connections to optimization theory. This kind of *principled approach for neural network design* is especially important for the broader signal processing community, since most existing neural network design principles from other communities are heuristic, without analogous counterparts based in theory. Finally, the proposed history-cognizant unrolling improved upon conventional unrolling for a number of different optimization algorithms, both quantitatively and visibly for various accelerated MRI scenarios. Importantly, this

improvement was achieved without a cost in the number of trainable parameters, in training time or in computational complexity.

The proposed history-cognizant unrolling methodology leads to several skip connections across iterations, leading to a densely-connected unrolled neural network. Dense neural networks have been shown to improve training by facilitating information flow through the network in forward propagation and by relieving issues concerned with gradient vanishing in backward propagation [49]. One of the advantages of these dense connections in the proposed unrolling is that the only computational overhead compared to conventional unrolling is the calculation of the 1×1 convolutions from the previous iterations. Thus, once trained, the computational time is similar to the conventional counterpart, as depicted in Table I.

One of the most commonly used accelerated PGD approaches is based on Nesterov acceleration [41]. In this case, the gradient descent is calculated at a linear combination of the previous two iterates weighted by factors α and $1 - \alpha$ respectively. Theoretically, this approach achieves the optimal convergence rate for conventional first-order optimization methods [41]. As such, Nesterov-type unrolling for DL reconstruction was readily explored in [23] for coded-illumination microscopy, where the unrolled accelerated PGD network used the history of the two previous outputs of the proximal units. Thus, we also performed a comparison between the proposed HC-PGD unrolling and Nesterov-type PGD unrolling [23], using the setup in Section III-B. Results in Figure 10 show that the proposed history-cognizant Dense-RNN architecture suppresses residual artifacts that are apparent in Nesterov-PGD unrolling, while also improving the quantitative metrics.

In our experiments, the history-cognizant unrolling visibly improved the performance of PGD-based methods compared to its conventional counterpart for all datasets. In the VS setting, especially for the coronal PD datasets, the improvement is less pronounced. Coronal PD datasets have higher SNR compared to PD-FS datasets [19]. In this higher SNR case, the conventional unrolling of the VS algorithm readily provides a high-quality reconstruction. Thus, additional improvement with history cognizant unrolling of VS is less apparent than in the PGD setting.

Our experiments in multi-coil MRI reconstruction predominantly utilized uniform undersampling patterns. While random undersampling patterns are commonly used in literature for training and testing DL-based reconstruction, uniform undersampling remains the most clinically used approach, and was considered to be a harder problem compared to random undersampling in [19, 60], as detailed in Section III-E. Uniform undersampling also creates a difficult problem for conventional non-DL methods that use a pre-specified regularizers, such as an ℓ_1 -wavelet regularization term. Figure 11 shows results from traditional non-DL implementations of PGD and ADMM with an ℓ_1 -wavelet regularization term. Even though the soft-thresholding parameter was manually fine-tuned for this slice for both PGD and ADMM, residual artifacts remain and quantitative metrics are noticeably worse than the DL methods. Both observations are consistent with previous studies [19]. Furthermore, these traditional methods require manual tuning of the hyperparameters, which vary across subjects and slices. We also note that we did not compare our methods to more

advanced non-DL methods [61–64], as our study focused on deep learning reconstruction. Furthermore, all our experiments concentrated on physics-driven DL reconstruction methods for MRI. While non-physics driven DL methods can speed up convergence and reduce the overall training time by learning the encoding matrix [65], this was not explored in this work which focused on algorithm unrolling for physics-based DL methods.

In this study, parameters of the proximal units were shared across unrolled iterations, consistent with the optimization literature where the regularization step remains unchanged over iterations. However, reconstruction quality may be further improved by learning different parameters, albeit at the cost of requiring more training examples and more careful training [27]. In addition, the training loss was calculated on the network output and reference target data which were both in image domain. While image-domain loss is a common practice in the literature, the loss defined over multi-coil k-space data may provide additional benefits for reducing artifacts [30, 31, 66], but this was not explored in the current study.

To the best of our knowledge, the Dense-RNN architecture that arises from the proposed history-cognizant unrolling have not been used for regression problems in general, and inverse problems in particular. However, Dense-RNNs have been independently proposed for different problems, such as scene parsing or labeling [67] where semantic information of an image segment is to be extracted. In such applications, RNNs are used for each image unit to receive the dependencies from other units through recurrent information forwarding between adjacent units, which decays for further units. In this setting, Dense-RNNs improve labeling by capturing long-range semantic dependencies among image units, highlighting a similar memory dependency as in this study.

Finally, the proposed history-cognizant unrolling combines past updates linearly. However, given the success of non-linear combination of the history of the gradients in improving the performance of gradient descent approaches [68], generalizing history cognizant unrolling to a non-linear scheme is an interesting extension of this work. However, this would require a theoretical underpinning for the use of non-linear combination of history gradients in solving the objective function in Equation (2), and thus is beyond the scope of this study.

V. Conclusion

We proposed a history-cognizant approach to unroll a given optimization algorithm for solving inverse problems in medical imaging using physics-driven deep learning. The proposed approach leads to an unrolled dense recurrent neural network architecture. From a network-design perspective, this consequent architecture enhances training and performance by facilitating information flow through the network. Accelerated MRI reconstruction experiments show that the proposed unrolling approach may considerably enhance reconstruction quality without substantially modifying computational complexity.

Acknowledgment

MRI data were obtained from the NYU fastMRI initiative database [50] which was acquired with the relevant institutional review board approvals as detailed in [50]. NYU fastMRI investigators provided data but did not

participate in analysis or writing of this report. A listing of NYU fastMRI investigators, subject to updates, can be found at fastmri.med.nyu.edu.

This work was supported by NIH P41EB027061, NIH U01EB025144, NSF CAREER CCF-1651825, NSF CMMI-1727757, CIF-1910385, ARO grant 73202-CS.

References

- [1]. Pruessmann KP, Weiger M, Bornert P, and Böesiger P, “Advances in sensitivity encoding with arbitrary k-space trajectories,” *Magn Reson Med*, vol. 46, no. 4, pp. 638–651, 2001. [PubMed: 11590639]
- [2]. Sutton BP, Noll DC, and Fessler JA, “Fast, iterative image reconstruction for MRI in the presence of field inhomogeneities,” *IEEE Trans Med Imaging*, vol. 22, no. 2, pp. 178–188, 2003. [PubMed: 12715994]
- [3]. Lustig M, Donoho D, and Pauly J, “Sparse MRI: The application of compressed sensing for rapid MR imaging,” *Magn Reson Med*, vol. 58, pp. 1182–95, 2007. [PubMed: 17969013]
- [4]. Block KT, Uecker M, and Frahm J, “Undersampled radial MRI with multiple coils. iterative image reconstruction using a total variation constraint,” *Magn Reson Med*, vol. 57, no. 6, pp. 1086–1098, 2007. [PubMed: 17534903]
- [5]. Akçakaya M, Nam S, et al., “Compressed sensing with wavelet domain dependencies for coronary MRI: a retrospective study,” *IEEE Trans Med Imaging*, vol. 30, no. 5, pp. 1090–1099, 2010.
- [6]. Knoll F, Bredies K, Pock T, and Stollberger R, “Second order total generalized variation (TGV) for MRI,” *Magn Reson Med*, vol. 65, no. 2, pp. 480–491, 2011. [PubMed: 21264937]
- [7]. Akçakaya M, Basha TA, et al., “Low-dimensional-structure self-learning and thresholding: regularization beyond compressed sensing for MRI reconstruction,” *Magn Reson Med*, vol. 66, no. 3, pp. 756–767, 2011. [PubMed: 21465542]
- [8]. Wang S, Su Z, et al., “Accelerating magnetic resonance imaging via deep learning,” in *Proc IEEE ISBI*, 2016, pp. 514–517.
- [9]. Kwon K, Kim D, and Park H, “A parallel MR imaging method using multilayer perceptron,” *Medical Physics*, vol. 44, no. 12, pp. 6209–6224, 2017. [PubMed: 28944971]
- [10]. Lee D, Yoo J, Tak S, and Ye JC, “Deep residual learning for accelerated MRI using magnitude and phase networks,” *IEEE Trans Biomed Eng*, vol. 65, no. 9, pp. 1985–1995, 2018. [PubMed: 29993390]
- [11]. Wang G, Ye JC, Mueller K, and Fessler JA, “Image reconstruction is a new frontier of machine learning,” *IEEE Trans Med Imaging*, vol. 37, pp. 1289–1296, 2018. [PubMed: 29870359]
- [12]. Dar SH and Çukur T, “Transfer learning for reconstruction of accelerated MRI acquisitions via neural networks,” Paris, France, 2018, *Proc ISMRM*.
- [13]. Chen H, Zhang Y, et al., “Low-dose CT with a residual encoder-decoder convolutional neural network,” *IEEE Trans Med Imaging*, vol. 36, pp. 2524–2535, 2017. [PubMed: 28622671]
- [14]. Han Y, Sunwoo L, and Ye JC, “k-space deep learning for accelerated MRI,” *IEEE Trans Med Imaging*, 2019.
- [15]. Han Y, Yoo J, et al., “Deep learning with domain adaptation for accelerated projection-reconstruction MR,” *Magn Reson Med*, vol. 80, pp. 1189–1205, 2018. [PubMed: 29399869]
- [16]. Yang Q, Yan P, et al., “Low-dose CT image denoising using a generative adversarial network with wasserstein distance and perceptual loss,” *IEEE Trans Med Imaging*, vol. 37, no. 6, pp. 1348–1357, 2018. [PubMed: 29870364]
- [17]. Kang E, Chang W, Yoo J, and Ye JC, “Deep convolutional framelet denosing for low-dose CT via wavelet residual network,” *IEEE Trans Med Imaging*, vol. 37, no. 6, pp. 1358–1369, 2018. [PubMed: 29870365]
- [18]. Wu Z, Sun Y, et al., “Simba: Scalable inversion in optical tomography using deep denoising priors,” preprint arXiv:1911.13241, 2019.
- [19]. Hammernik K, Klatzer T, et al., “Learning a variational network for reconstruction of accelerated MRI data,” *Magn Reson Med*, vol. 79, pp. 3055–3071, 2018. [PubMed: 29115689]

- [20]. Schlemper J, Caballero J, Hajnal JV, Price AN, and Rueckert D, "A deep cascade of convolutional neural networks for dynamic MR image reconstruction," *IEEE Trans Med Imaging*, vol. 37, no. 2, pp. 491–503, 2017. [PubMed: 29035212]
- [21]. Mardani M, Monajemi H, et al., "Recurrent generative adversarial networks for proximal learning and automated compressive image recovery," preprint arXiv:1711.10046, 2017.
- [22]. Zhang J and Ghanem B, "ISTA-Net: Interpretable optimization-inspired deep network for image compressive sensing," in *Proc IEEE CVPR*, 2018, pp. 1828–1837.
- [23]. Kellman M, Bostan E, Repina N, and Waller L, "Physics-based learned design: Optimized coded-illumination for quantitative phase imaging," *IEEE Trans Comp Imaging*, 2019.
- [24]. Yang Y, Sun J, Li H, and Xu Z, "ADMM-CSNet: A deep learning approach for image compressive sensing," *IEEE Trans Pat Anal Mach Intel*, 2018.
- [25]. Yang Y, Sun J, Li H, and Xu Z, "Deep ADMM-Net for compressive sensing MRI," in *Advances in neural information processing systems*, 2016, pp. 10–18.
- [26]. Ye DH, Srivastava S, Thibault J-B, Sauer K, and Bouman C, "Deep residual learning for model-based iterative CT reconstruction using plug-and-play framework," in *Proc IEEE ICASSP*, 2018, pp. 6668–6672.
- [27]. Aggarwal HK, Mani MP, and Jacob M, "Modl: Model-based deep learning architecture for inverse problems," *IEEE Trans Med Imaging*, vol. 38, no. 2, pp. 394–405, 2018. [PubMed: 30106719]
- [28]. Qin C, Schlemper J, et al., "Convolutional recurrent neural networks for dynamic MR image reconstruction," *IEEE Trans Med Imaging*, vol. 38, pp. 280–290, 2018. [PubMed: 30080145]
- [29]. Duan J, Schlemper J, et al., "VS-Net: Variable splitting network for accelerated parallel MRI reconstruction," in *Proc MICCAI*, 2019, pp. 713–722.
- [30]. Yaman B, Hosseini SAH, et al., "Self-supervised physics-based deep learning MRI reconstruction without fully-sampled data," *Proc IEEE ISBI*, 2020.
- [31]. Yaman B, Hosseini SAH, et al., "Self-supervised learning of physics-guided reconstruction neural networks without fully-sampled reference data," *Magn Reson Med*, 2020.
- [32]. Adler J and Oktem O, "Learned Primal-Dual Reconstruction," *IEEE Trans Med Imaging*, vol. 37, pp. 1322–1332, 2018. [PubMed: 29870362]
- [33]. Cheng J, Wang H, Ying L, and Liang D, "Model learning: Primal dual networks for fast MR imaging," in *Proc MICCAI*, 2019, pp. 21–29.
- [34]. Bostan E, Kamilov US, and Waller L, "Learning-based image reconstruction via parallel proximal algorithm," *IEEE Sig Proc Let*, vol. 25, pp. 989–993, 2018.
- [35]. Kamilov US and Mansour H, "Learning optimal non-linearities for iterative thresholding algorithms," *IEEE Sig Proc Letters*, vol. 23, no. 5, pp. 747–751, 2016.
- [36]. Akçakaya M, Moeller S, Weingärtner S, and Urbil K, "Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging," *Magn Reson Med*, vol. 81, no. 1, pp. 439–453, 2019. [PubMed: 30277269]
- [37]. Hosseini SAH, Zhang C, et al., "Accelerated coronary mri with sraki: A database-free self-consistent neural network k-space reconstruction for arbitrary undersampling," *PLOS One*, vol. 15, no. 2, pp. e0229418, 2020. [PubMed: 32084235]
- [38]. Hosseini SAH, Zhang C, Uğurbil K, Moeller S, and Akçakaya M, "sRAKI-RNN: accelerated MRI with scan-specific recurrent neural networks using densely connected blocks," in *SPIE Wavelets and Sparsity XVIII*, 2019, p. 111381B.
- [39]. Knoll F, Hammernik K, et al., "Deep-learning methods for parallel magnetic resonance imaging reconstruction: A survey of the current approaches, trends, and issues," *IEEE Sig Proc Mag*, vol. 37, no. 1, pp. 128–140, 2020.
- [40]. Nesterov Y, "Smooth minimization of non-smooth functions," *Mathematical Programming*, vol. 103, no. 1, pp. 127–152, 2005.
- [41]. Nesterov Y, "Gradient methods for minimizing composite functions," *Mathematical Programming*, vol. 140, no. 1, pp. 125–161, 2013.
- [42]. Nesterov Y and Florea M, "Gradient methods with memory," *Tech. Rep*, 2019.

- [43]. Beck A and Teboulle M, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J Imag Sci*, vol. 2, no. 1, pp. 183–202, 2009.
- [44]. Chambolle A and Pock T, "A remark on accelerated block coordinate descent for computing the proximity operators of sum of convex functions," *Tech. Rep*, 2015.
- [45]. Hong M, Wang X, Razaviyayn M, and Luo Z-Q, "Iteration complexity analysis of block coordinate descent methods," *Mathematical Programming Series A*, 2016.
- [46]. Tseng P, "On accelerated proximal gradient methods for convex-concave optimization," 2008, preprint.
- [47]. Becker S, Bobin J, and Candès EJ, "NESTA: A fast and accurate first-order method for sparse recovery," *SIAM J Imag Sci*, vol. 4, pp. 1–39, 2011.
- [48]. He K, Zhang X, Ren S, and Sun J, "Deep residual learning for image recognition," in *Proc IEEE CVPR*, 2016, pp. 770–778.
- [49]. Huang G, Liu Z, Van Der Maaten L, and Weinberger KQ, "Densely connected convolutional networks," in *Proc IEEE CVPR*, 2017, pp. 4700–4708.
- [50]. Zbontar J, Knoll F, et al., "fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning," *Radiology: AI*, vol. 2, no. 1, 2020.
- [51]. Fessler JA, "Optimization methods for magnetic resonance image reconstruction: Key models and optimization algorithms," *IEEE Signal Processing Magazine*, vol. 37, no. 1, pp. 33–40, 2020. [PubMed: 32317844]
- [52]. Combettes PL and Pesquet J-C, "Proximal splitting methods in signal processing," *arXiv:0912.3522*, 2009.
- [53]. Goldstein T and Osher S, "The split bregman method for L1-regularized problems," *SIAM journal on imaging sciences*, vol. 2, no. 2, pp. 323–343, 2009.
- [54]. Ramani S and Fessler JA, "Parallel MR image reconstruction using augmented lagrangian methods," *IEEE Trans Med Imaging*, vol. 30, no. 3, pp. 694–706, 2010. [PubMed: 21095861]
- [55]. Boyd S, Parikh N, et al., "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found T Mach Learn*, pp. 1–122, 2011.
- [56]. Chambolle A and Pock T, "A first-order primal-dual algorithm for convex problems with applications to imaging," *J Math Imag Vis*, vol. 40, pp. 120–145, 2011.
- [57]. Sotiropoulos SN, Moeller S, et al., "Effects of image reconstruction on fiber orientation mapping from multichannel diffusion MRI: reducing the noise floor using SENSE," *Magn Reson Med*, pp. 1682–1689, 2013. [PubMed: 23401137]
- [58]. Uecker M, Lai P, et al., "ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA," *Magn Reson Med*, vol. 71, no. 3, pp. 990–1001, 2014. [PubMed: 23649942]
- [59]. Timofte R, Agustsson E, Van Gool L, Yang M-H, and Zhang L, "Ntire 2017 challenge on single image super-resolution: Methods and results," in *Proc IEEE CVPR*, 2017, pp. 114–125.
- [60]. Knoll F, Hammernik K, et al., "Assessment of the generalization of learned image reconstruction and the potential for transfer learning," *Magn Reson Med*, vol. 81, pp. 116–128, 2019. [PubMed: 29774597]
- [61]. Qu X, Guo D, et al., "Undersampled MRI reconstruction with patch-based directional wavelets," *Magn Res Imag*, vol. 30, no. 7, pp. 964–977, 2012.
- [62]. Zhan Z, Cai J-F, et al., "Fast multiclass dictionaries learning with geometrical directions in MRI reconstruction," *IEEE Trans Bio Eng*, pp. 1850–1861, 2015.
- [63]. Qu X, Hou Y, et al., "Magnetic resonance image reconstruction from undersampled measurements using a patch-based nonlocal operator," *Med Image Anal*, vol. 18, no. 6, pp. 843–856, 2014. [PubMed: 24176973]
- [64]. Zhang X, Guo D, et al., "Image reconstruction with low-rankness and self-consistency of k-space data in parallel MRI," *Med Image Anal*, p. 101687, 2020. [PubMed: 32294605]
- [65]. Giryes R, Eldar YC, Bronstein AM, and Sapiro G, "Tradeoffs between convergence speed and reconstruction accuracy in inverse problems," *IEEE Trans Sig Proc*, vol. 66, no. 7, pp. 1676–1690, 2018.

- [66]. Tamir JI, Yu SX, and Lustig M, “Unsupervised deep basis pursuit: Learning inverse problems without ground-truth data,” preprint arXiv:1910.13110, 2019.
- [67]. Fan H, Chu P, Latecki LJ, and Ling H, “Scene parsing via dense recurrent neural networks with attentional selection,” in Proc IEEE WACV, 2019, pp. 1816–1825.
- [68]. Kingma DP and Ba J, “Adam: A method for stochastic optimization,” arXiv:1412.6980, 2014.

Fig. 1:

(a) The conventional unrolling for iterative algorithms for inverse problems and (b) the proposed history-cognizant unrolling leading to a dense recurrent neural network architecture. In both cases, the algorithms are unrolled for T iterations, where each unrolled iteration consists of proximal (or regularization) unit (P), and a data consistency unit (DC), with $x^{(0)} = A^H y$ as input to the network. The P unit is implemented implicitly via a neural network, while the DC unit corresponds to a gradient descent step for the proximal gradient descent algorithms and typically to conjugate gradient for variable splitting methods. For the proposed history-cognizant unrolling, an additional 1×1 convolutional layer combines all previous proximal operator outputs prior to the DC unit, where C denotes concatenation.

Fig. 2:

Summary of the training procedure in the proposed history-cognizant unrolling of inverse problem optimization algorithms with a Dense-RNN architecture for accelerated DL MRI reconstruction. Parameters of the unrolled network, which are dominantly characterized by proximal (P) operator units, are updated in each epoch based on a given loss metric using conventional end-to-end training. Data consistency (DC) units utilize the acquired data, coil sensitivity maps and undersampling mask in updates. The uniform undersampling mask utilized for training the network is also depicted.

Fig. 3:

A representative slice from the Coronal PD dataset with 4-fold acceleration rate using a uniform undersampling pattern, reconstructed with DL-MRI reconstruction based on the conventional PGD and the proposed HC-PGD unrolling, as well as error images with respect to reference fully-sampled image scaled by a factor of 10. Zero-filled image corresponds to the input to the network, $x^{(0)}$. The green rectangle in the reference image marks the zoomed area shown in the second row for each method. Residual aliasing artifacts, visible in PGD unrolling results, have been suppressed using the HC-PGD scheme. Improvements in the

quantitative metrics, shown on the lower-left corner of the images, also align with these visual observations.

Fig. 4:

A slice from the Coronal PD-FS dataset with 4-fold acceleration using uniform undersampling, reconstructed with DL-MRI based on conventional PGD and proposed HC-PGD unrolling methods, and error images (scaled by 10). The green rectangle in the reference image mark the zoom area. The proposed HC-PGD unrolling successfully removes some residual artifacts that are visible in the PGD, and improves the quantitative metrics.

Fig. 5:

Summary of the quantitative metrics for reconstruction quality, (a) PSNR and (b) SSIM in Coronal PD and Coronal PD-FS datasets using the conventional PGD and the proposed HC-PGD methods. The boxes mark the interquartile range and median of the metrics. The proposed history-cognizant unrolling of the PGD has visibly enhanced reconstruction, based on both metrics and in both datasets. Statistical analysis based on Wilcoxon signed rank test further confirm that all improvements are statistically significant.

Fig. 6:

Results of changing the proximal operator CNN, with a slice from the Coronal PD dataset with 4-fold uniform undersampling, reconstructed with DL-MRI reconstruction using the PGD and HC-PGD methods unrolled in a network with a U-Net design for the proximal operator units, along with Error images ($\times 5$), with the green rectangle marking the zoom-up area. Similar to the ResNet experiments, the proposed HC-PGD outperforms PGD in suppressing visible residual artifacts.

Fig. 7:

Results on random undersampling experiments from (a) the Coronal PD and (b) Coronal PD-FS datasets, at a 4-fold acceleration rate, reconstructed with DL-MRI reconstruction based on the conventional PGD and the proposed HC-PGD unrolling, as well as error images ($\times 10$), with the green rectangle showing the zoom-up area. Regardless of the undersampling pattern, the history-cognizant unrolling approach, HC-PGD suppresses the residual aliasing visible in the conventional PGD unrolling, as marked by the arrows. The quantitative metrics align with these observations.

Fig. 8:

Example slice from the Coronal PD with 4-fold uniform undersampling using DL-MRI reconstruction based on the conventional VSQP and ADMM, and the proposed HC-VSQP and HC-ADMM unrolling, as well as error images ($\times 10$), with the green rectangle showing the zoomed area. The history-cognizant unrolling approaches, HC-VSQP and HC-ADMM suppress the residual aliasing visible in the conventional VSQP and ADMM unrolling, marked by the arrows. Improvements are also observed in the quantitative metrics.

Fig. 9:

Example slice from the Coronal PD-FS with 4-fold uniform undersampling using DL-MRI reconstruction based on the conventional VSQP and ADMM, and the proposed HC-VSQP and HC-ADMM unrolling, as well as error images ($\times 10$), with green rectangle showing the zoomed area. The reconstructions from the history-cognizant unrolling approach, HC-VSQP and HC-ADMM suppress the residual aliasing visible in the conventional VSQP and ADMM unrolling, marked by the arrows.

Fig. 10:

The effect of Nesterov-type unrolling versus full history cognizant unrolling in HC-PGD for a slice from the Coronal PD dataset with 4-fold uniform undersampling, along with Error images ($\times 5$), with the green rectangle marking the zoom-up area. Aliasing artifacts remain in the reconstructed images from conventional PGD and Nesterov-type unrolling, while these are suppressed with the proposed HC-PGD unrolling.

Fig. 11:

Comparison of traditional non-DL PGD and ADMM, both using ℓ_1 -wavelet regularization, against DL methods based on HC-PGD and HC-ADMM unrolling for a Coronal PD slice with 4-fold uniform undersampling, along with error images ($\times 10$), with the green rectangle marking the zoom-up area. Aliasing artifacts remain in the reconstructed images from traditional PGD and ADMM, while DL versions of these algorithms successfully suppress these artifacts, resulting in superior reconstruction quality and quantitative metrics, consistent with previous studies [19].

TABLE I:

Average and standard deviation (in ms) of reconstruction time over all test slices for PGD, HC-PGD, VSQP, HC-VSQP, ADMM and HC-ADMM in Coronal PD and Coronal PD-FS datasets. History cognizant unrolling has very slight impact on reconstruction time compared to its conventional counterpart. Furthermore, VSQP and ADMM methods with CG-based data consistency units, are considerably slower than PGD methods.

	PGD	HC-PGD	VSQP	HC-VSQP	ADMM	HC-ADMM
Coronal PD	197 $\pm$ 4.2 ms	197 $\pm$ 7.2 ms	340 $\pm$ 6.8 ms	325 $\pm$ 6.1 ms	316 $\pm$ 5.0 ms	320 $\pm$ 6.0 ms
Coronal PD-FS	197 $\pm$ 5.6 ms	198 $\pm$ 6.6 ms	339 $\pm$ 5.6 ms	325 $\pm$ 6.2 ms	321 $\pm$ 6.1 ms	318 $\pm$ 6.1 ms