An adaptive intelligence algorithm for under sampled knee MRI reconstruction

Pezzotti, N.; Yousefi, S.; Elmahdy, M.S.; Gemert, J.H.F. van; Schuelke, C.; Doneva, M.; ... ; Staring, M.

Note: To cite this publication please use the final published version (if applicable).

An Adaptive Intelligence Algorithm for Undersampled Knee MRI Reconstruction

NICOLA PEZZOTTI[1], SAHAR YOUSEFI[2], MOHAMED S. ELMAHDY[2], JEROEN HENDRIKUS FRANSISCUS VAN GEMERT[3], CHRISTOPHE SCHUELKE[4], MARIYA DONEVA[4], (Member, IEEE), TIM NIELSEN[4], SERGEY KASTRYULIN[5], BOUDEWIJN P. F. LELIEVELDT[2,6], (Member, IEEE), MATTHIAS J. P. VAN OSCH[2], ELWIN DE WEERDT[3], AND MARIUS STARING[2,6]

[1]Philips Research, 5656AE Eindhoven, The Netherlands
[2]Department of Radiology, Leiden University Medical Center, 2333 ZA Leiden, The Netherlands
[3]Philips Healthcare, 5684PC Best, The Netherlands
[4]Philips Research, 22335 Hamburg, Germany
[5]Philips Research, 121205 Moscow, Russia
[6]Department of Pattern Recognition and Bioinformatics, Delft University of Technology, 628CD Delft, The Netherlands

Corresponding author: Nicola Pezzotti (nicola.pezzotti@philips.com)

Elwin de Weerdt and Marius Staring are co-senior authors.

ABSTRACT Adaptive intelligence aims at empowering machine learning techniques with the additional use of domain knowledge. In this work, we present the application of adaptive intelligence to accelerate MR acquisition. Starting from undersampled k-space data, an iterative learning-based reconstruction scheme inspired by compressed sensing theory is used to reconstruct the images. We developed a novel deep neural network to refine and correct prior reconstruction assumptions given the training data. The network was trained and tested on a knee MRI dataset from the 2019 fastMRI challenge organized by Facebook AI Research and NYU Langone Health. All submissions to the challenge were initially ranked based on similarity with a known groundtruth, after which the top 4 submissions were evaluated radiologically. Our method was evaluated by the fastMRI organizers on an independent challenge dataset. It ranked #1, shared #1, and #3 on respectively the 8× accelerated multi-coil, the 4× multi-coil, and the 4× single-coil tracks. This demonstrates the superior performance and wide applicability of the method.

INDEX TERMS Image reconstruction, MRI, deep learning, ISTA, fastMRI challenge.

I. INTRODUCTION

Magnetic Resonance Imaging (MRI) is a widely applied non-invasive imaging modality, with excellent soft tissue contrast and high spatial resolution. Unlike Computed Tomography (CT) scanning, MRI does not expose patients to any ionizing radiation, making it a compelling alternative. MR images are essential for clinical assessment of soft tissue as well as functional and structural measurements, which leads to early detection and diagnosis of many diseases. However, MRI is relatively slow compared to other imaging modalities. The total examination time can vary from 15 minutes for knee imaging to an hour or more for cardiac imaging. Remaining still for this long in a confined space is challenging for any patient, being especially difficult for children, elderly and patients under pain. Motion artifacts are not only difficult to

The associate editor coordinating the review of this manuscript and approving it for publication was Pengcheng Liu.

correct, which may require a complete re-scan [1]. Furthermore, the acquisition time affects the temporal resolution and subsequently limits the potential of MRI for dynamic imaging, where high temporal resolution and robustness against motion are critical for diagnosis. Moreover, the relatively long scan times lead to high costs that limit the availability of MRI scanners [2]. Therefore, fast acquisition and reconstruction are crucial to improve the performance of current MR scanners, which led in recent years to the development of techniques such as parallel reception, compressed sensing and multi-band accelerations. However, there is still a need for further scan acceleration.

The long acquisition time is intrinsic to the scanner and physics properties of MRI. For the majority of scans performed in clinical practice, this acquisition is done through consecutive reading-out of single lines in k-space. These readouts are constrained by physical limitations of the hardware, the contrast generating principle, and

human physiology. The scanning time could be shortened by reducing the number of acquired lines in k-space, i.e. by undersampling the 2D or 3D k-space. However, this could violate the Nyquist criterion, resulting in aliasing and blurriness in the reconstructed images, rendering them unqualified for clinical purpose. Compressed Sensing (CS) and Parallel Imaging are the most common solutions for acceleration by undersampling, while maintaining image quality. Compressed Sensing, the focus of this paper, introduced by Donoho [3], Lustig [4] and Candes [5], leverages the fact that MR images can be compressed in some domain, restoring the missing k-space data through an iterative reconstruction algorithm [6]. Parallel Imaging uses multiple receive coils that provide an additional signal encoding mechanism, allowing to reduce the number of necessary k-space lines to reconstruct an image, thus partially parallelizing the data acquisition [7].

When CS is used to accelerate MR acquisitions, the k-spaces is sampled pseudo-randomly and the image is subsequently reconstructed by promoting a sparse solution. In the optimal setting, the reconstructed image will be identical to the Fourier transform of the full k-space and have a limited number of large coefficients when transformed to the sparse domain. Equation (1) shows the optimization function that describes the CS algorithm:

$$\min_{\mathbf{x}} \left\{ \| M\mathcal{F}\mathbf{x} - M\mathbf{y} \|_2^2 + \lambda \| \Psi\mathbf{x} \|_1 \right\}, \quad (1)$$

where $\mathbf{x}$ is the reconstructed image, $\mathbf{y}$ is the fully measured k-space data, $\mathcal{F}$ is the Fourier transform, M (mask) is the undersampling operation, $\Psi\mathbf{x}$ represents the sparsity transform coefficients, and λ is the regularization parameter. The ℓ_1 norm is used to enforce sparsity of the solution in a domain specified by the transformation Ψ. The ℓ_2 norm is used as a similarity measure between the measured k-space data $M\mathbf{y}$ and the reconstructed k-space $M\mathcal{F}\mathbf{x}$, called the "data consistency" term. Note that, in case of multi-coil acquisitions, the data consistency term is given by:

$$\sum_q \| M\mathcal{F} \left(\mathbf{S}_q \cdot \mathbf{x} \right) - M\mathbf{y}_q \|_2^2, \quad (2)$$

where q denotes the coil element and $\mathbf{S_q}$ the corresponding coil sensitivity map. The coil sensitivity maps $\mathbf{S}$ are computed using the fully centered region of k-space. A low-passed version of the coil images $\mathbf{x}_q^{\text{lpf}}$ is obtained by cropping the available region of k-space. The sensitivity map $\mathbf{S}_q$, for the individual coil element is computed as follows:

$$\mathbf{S}_q = \frac{\mathbf{x}_q^{\text{lpf}}}{\sqrt{\sum_j \left(\mathbf{x}_j^{\text{lpf}} \right)^2}} \quad (3)$$

To simplify notation, without loss of generality, the single-coil data consistency term will be used throughout this paper.

Recently, deep learning has shown promising results for speeding up MR acquisition by adopting Convolutional Neural Networks (CNN) and Generative Adversarial Networks

(GAN). In contrast to iteratively solving optimization problems, deep learning offers a solution for reconstructing highly-accelerated scans by adopting learnable reconstruction schemes.

The literature of deep learning-based reconstruction algorithms can be divided into two categories [8]. First, data-driven approaches, where a neural network is trained to find the optimal transformation from the zero-filled k-space to the desired reconstruction. Here, the network is completely dependent on the underlying training dataset without any task-specific prior knowledge on the domain; following are selected exemplar algorithms of this approach. Quan *et al.* [9] developed a GAN network for MR reconstruction starting from undersampled data. Their network consists of two consecutive networks, one for reconstruction and one for refining the results. They used a cyclic data consistency term alongside the WGAN loss. Mardani *et al.* [10] developed a GAN network for CS. The proposed network corrects aliasing artifacts of MR images. Guo *et al.* [11] proposed a WGAN with recurrent context-awareness to reconstruct MRI images from highly undersampled k-space data. Schlemper *et al.* propose a cascaded CNN-based compressive sensing (CS) technique for the reconstruction of diffusion tensor cardiac MRI [12]. Yang *et al.* proposed a conditional GAN-based architecture for de-aliasing and fast CS-MRI [13], [14]. Putzky *et al.* [15] treated the MR reconstruction problem as an inverse problem. They applied the previously introduced invertible Recurrent Inference Machine (i-RIM) model [16], which iteratively updates its current state based on the output of the forward model. The model was trained and evaluated on the single- and multi-coil data at $4\times$ and $8\times$ accelerations from the fastMRI challenge (see Section II for more details). AUTOMAP [17] reports good reconstruction results with an architecture that learns to directly transform k-space into image data. Lee *et al.* [18] introduced two separate deep residual networks for magnitude and phase. The proposed networks successfully reconstructed images even when obtained with high undersampling factors.

Second, hybrid approaches are presented in the literature. This class of algorithms builds on top of existing reconstruction solutions and integrate learning-based approaches to substitute part of the original computations, often by adopting an unrolled implementation of an iterative algorithm [19]. A notable example is the Variational network presented by Hammernik *et al.* [20] utilizing learned filters in an existing iterative optimization scheme, while Yang *et al.* presented the Deep ADMM-Net [21], which extends the Alternating Direction Method of Multipliers (ADMM) [22] approach by integrating learnable operators.

Aggarwal *et al.* [23] introduced a model based deep learning architecture named MoDL to solve the inverse problem, including MR reconstruction. The proposed model consists of a series of recursive linear CNN networks. These networks share weights for regularization and reduction in the number of parameters. The proposed network imitates the CS algorithm and for numerical optimization, the authors introduced

a data consistency term using a conjugate gradient (CG) optimization scheme at every iteration. The model was trained on multi-coil brain MR slices from 4 patients and tested on one patient. Ramzi *et al.* [24] provided a reproducible benchmark of deep learning based reconstruction methods on the single-coil part of the fastMRI dataset [25]. The benchmark consists of a U-net [26], cascade net [27], KIKI-net [28], and PD-net [29]. Cascade net has been inspired by a dictionary learning approach [30]. This approach is composed of residual convolutional blocks applied in image space followed by data consistency layers. The data consistency layers enforce the k-space values be close to the original k-space measurements. KIKI-net is a cascaded network where a non-residual convolutional block has been added to perform k-space completion, while PD-net provides a learnable and unrolled version of the Primal Dual Hybrid Gradient optimization algorithm [31]. Seitzer *et al.* discussed the inadequacy of loss function for training a CS-MRI reconstruction CNN [32]. In that study they proposed a refinement method which incorporates both loss functions in a harmonious way to improve the training stability.

Recently, Zhang and Ghanem [33] developed a deep learning approach called ISTA-Net that mimics the conventional ISTA algorithm, but enriches it by replacing the sparsifying transform and the thresholding with learned operations. The resulting network does not implement a fully iterative algorithm, but it simulates it by adopting a fixed number of iterations, effectively enabling the implementation of a deep neural network that can be trained by the backpropagation algorithm. Inspired by the work of Zhang and Ghanem [33], in this paper we propose a deep-learning based solution, Adaptive-CS-Network, that mimics the ISTA algorithm, but introduces strong prior information, i.e., inductive biases, to better constrain the reconstruction problem. The main contributions of this work are: i) we propose a novel CNN network that integrates and enhances the conventional CS approach; ii) it integrates multiscale sparsification, inspired by wavelet transforms, but in a learnable manner; iii) we adopt domain-specific knowledge, such as data consistency, a prior on known phase behavior, and the location of the background: these computations cannot be easily learned by a CNN; iv) the proposed model exploits the correlation between neighbouring slices by adopting a 2.5D learning approach. In addition, we propose a hierarchical training strategy that leverages the available data. We conducted extensive experiments to investigate the performance of the network, and show that domain specific information is crucial for reconstructing high-quality MR images. The proposed network showed superior performance by winning one, and co-winning a second track out of the three tracks of the fastMRI challenge [25].

II. FastMRI CHALLENGE

The fastMRI challenge is a challenge organized by Facebook AI Research and NYU Langone Health [25]. The aim of the challenge is to advance and encourage AI-based research in MR reconstruction in order to allow acceleration of the acquisition and, subsequently, to reduce the examination time. The challenge is divided in three tracks: 4× single-coil, 4× multi-coil, and 8× multi-coil accelerations. Eight teams participated in the multi-coil track and 17 teams in the single-coil track [34].

A. DATASET

The challenge organizers released a large-scale dataset of raw MR data of the knee [35]. The data was acquired with a 2D protocol in the coronal direction with a 15 channel knee coil array using Siemens MR machines at two different field strengths: 1.5T and 3T [25]. The data was acquired using two pulse sequences: a *proton density weighting* with (PDFS) and without (PD) fat suppression. The data is divided approximately equally between these pulse sequences. The pixel size is 0.5 mm × 0.5 mm with a slice thickness of 3 mm.

The dataset is divided in 4 categories: training (973 volumes, 34,742 slices), validation (199 volumes, 7,135 slices), test (118 volumes, 4,092 slices), and challenge (104 volumes, 3,810 slices). These numbers are the same for multi-coil and single-coil data, with the exception of the test and challenge categories, where single-coil data has respectively 10 and 12 volumes less than the multi-coil data. The training, validation and test sets were publicly available since late November, 2018, while the challenge set was available since September 2019. The full k-space was available for all the datasets except for the test and challenge sets. Training and validation sets were considered for training and optimizing our model, while the test set was used for evaluating model performance on a public leaderboard. The final model was evaluated by the organizers on the independent challenge set.

The k-space data provided in the challenge were undersampled using a Cartesian mask, where k-space lines are set to zero in the phase encoding direction. The sampling density is dependent on the acceleration rate (4× or 8×), where the sampled lines are randomly selected. All masks, however, are fully sampled in the central area of k-space which corresponds to the low frequencies of the image. For the 4× accelerated scans, this percentage is 8% while it is 4% for 8× acceleration. Besides making the reconstruction problem easier to solve, such lines allow for obtaining a low-pass filtered version of the image that is used to compute the coil sensitivity maps S_q as presented in Equation (3) using a root sum of square approach [25].

B. QUANTITATIVE EVALUATION

In order to measure the accuracy of the reconstructed volumes r compared to the target volumes t, the following metrics were considered:

1) NORMALIZED MEAN SQUARE ERROR (NMSE)

measures the square of the Euclidean norm between a pair of images:

$$\text{NMSE} = \frac{||\mathbf{r} - \mathbf{t}||_2^2}{||\mathbf{t}||_2^2} \tag{4}$$

2) PEAK SIGNAL-TO-NOISE RATIO (PSNR)

the ratio between the maximum intensity and the underlying distortion noise:

$$\text{PSNR} = 10 \log_{10} \frac{\max(\mathbf{t})^2}{\frac{1}{N}||\mathbf{r} - \mathbf{t}||_2^2} \tag{5}$$

3) STRUCTURAL SIMILARITY INDEX METRIC (SSIM)

measures image similarity using human perception aspects [36]. SSIM is calculated by measuring three image distortions including luminance $l(\cdot)$, contrast $c(\cdot)$ and structure $s(\cdot)$:

$$\text{SSIM} = l(\mathbf{r}, \mathbf{t})^\alpha c(\mathbf{r}, \mathbf{t})^\beta s(\mathbf{r}, \mathbf{t})^\gamma, \tag{6}$$

where α, β, γ are the distortion weights, here chosen as 1. In this study, similar to the fastMRI challenge, the SSIM score is computed on the magnitude version of the 2D MR scans, leading to grayscale images.

C. RADIOLOGICAL EVALUATION ON THE CHALLENGE DATASET

We submitted the reconstructions on the challenge dataset via an online form, which were then evaluated independently by the fastMRI organizers, described in detail by Knoll et al. [34]. All submissions were ranked by the SSIM metric, after which only the 4 highest ranking submissions were evaluated by a panel of 7 radiologists. The panel was asked to evaluate the reconstructions on a scale from 1 to 5 on four different categories, where 1 is the best and 5 is the worst. The 4 categories were the rating of artifacts, reconstruction sharpness, perceived contrast-to-noise ratio and diagnostic confidence. The radiological scores were subsequently averaged and translated to a final ranking.

III. METHODS

In this section we present the background of our solution, first by introducing the Iterative Shrinkage-Thresholding Algorithm (ISTA) [37] and, second, by introducing its deep learning-based variant, ISTA-Net [33]. Then, we present our solution, the Adaptive-CS-Network, that builds on top of the ISTA-Net framework by introducing several improvements, including strong inductive biases derived from domain knowledge on the reconstruction problem.

A. ISTA BACKGROUND

ISTA is an optimization algorithm to solve (1) in an iterative fashion, starting from the reconstruction $\mathbf{x}_0$, which is often obtained by reconstructing the zero-filled undersampled k-space. The initial estimate is refined using the following update rules:

$$\mathbf{r}_{i+1} = \mathbf{x}_i - \rho \mathcal{F}^T (M\mathcal{F}\mathbf{x}_i - M\mathbf{y}), \tag{7}$$

$$\mathbf{x}_{i+1} = \arg \min_{\mathbf{x}} \frac{1}{2} \| \mathbf{x} - \mathbf{r}_{i+1} \|_2^2 + \lambda \| \Psi\mathbf{x} \|_1, \tag{8}$$

where F^T denotes inverse Fourier transform, $\mathbf{r}_{i+1}$ is an update of the estimate $\mathbf{x}_i$, where the error in the measured

data $M\mathbf{y}$ is corrected by a step ρ. Equation (8) is a special case of the proximal mapping, with a regularization weight λ, and a crucial step for optimization algorithms such as ISTA, ADMM [22] and AMP [38]. When Ψ is a wavelet transform $\mathbf{W}$, it can be proven that

$$\mathbf{x}_{i+1} = \mathbf{W}^{-1} soft(\mathbf{W}\mathbf{r}_{i+1}, \lambda), \tag{9}$$

where $soft$ is the soft-tresholding operator defined as $soft(\mathbf{u}, \lambda) = \max(|\mathbf{u}| - \lambda, 0) \cdot \frac{\mathbf{u}}{|\mathbf{u}|}$. In general, solving (8) is not straightforward for non-linear operators Ψ, limiting the applicability of the ISTA framework to simple transforms. Another problem of this family of algorithms, is the difficulty of tuning the hyperparameters λ and ρ in addition to its slow convergence, hence requiring a lot of iterations to achieve the optimal solution of (1).

B. ISTA-NET

Recently, Zhang and Ghanem introduced a deep-learning approach to overcome the limitations of the ISTA framework for image-to-image reconstruction. Their solution, called ISTA-Net [33], replaces the handcrafted transform Ψ with a learned operator $\mathcal{S}(\cdot)$, which consists of a 2D learnable convolution followed by a rectified linear unit (ReLU) and a second convolution. By replacing Ψ with $\mathcal{S}(\cdot)$ in (8), we can rewrite the update rule as

$$\mathbf{x}_{i+1} = \arg \min_{\mathbf{x}} \frac{1}{2} \| \mathbf{x} - \mathbf{r}_{i+1} \|_2^2 + \lambda \| \mathcal{S}(\mathbf{x}) \|_1, \tag{10}$$

and, by defining $\hat{\mathcal{S}}$ as the inverse of $\mathcal{S}$, i.e., $\hat{\mathcal{S}} \circ \mathcal{S} = I$, Zhang and Ghanem propose to update (9) as follows:

$$\mathbf{x}_{i+1} = \hat{\mathcal{S}}(soft(\mathcal{S}(\mathbf{r}_{i+1}), \lambda)), \tag{11}$$

where $\hat{\mathcal{S}}$ has a similar architecture as $\mathcal{S}$.

The model is trained end-to-end, where the iterations of the ISTA algorithm are "unrolled", i.e., a number b of identical reconstruction blocks are created. Note that in the ISTA-Net approach, the learnable parameters are shared among all the blocks in the unrolled network, unlike our solution. The training loss is defined as a combination of the reconstruction and discrepancy loss:

$$\mathcal{L} = \mathcal{L}_{reconstruction} + \sigma \mathcal{L}_{discrepancy} \tag{12}$$

$$\mathcal{L}_{reconstruction} = \| \mathbf{x}_b - \mathcal{F}^T \mathbf{y} \|_2^2 \tag{13}$$

$$\mathcal{L}_{discrepancy} = \frac{1}{b} \sum_{i=1}^{b} \| \hat{\mathcal{S}}(\mathcal{S}(\mathbf{x}_i)) - \mathbf{x}_i \|_2^2 \tag{14}$$

The reconstruction loss encodes the need for the final reconstruction, defined as $\mathbf{x}_b$, to be as close as possible in the least squares sense to the ground-truth image, i.e., $F^T\mathbf{y}$. The discrepancy loss stimulates that $\hat{\mathcal{S}} \circ \mathcal{S} = I$. The σ parameter allows to control the weight given to the discrepancy loss, and it is chosen to be arbitrarily small, e.g., $\sigma = 0.01$. An extension, called ISTA-Net$^+$ is also presented by the authors, where residual computations are adopted.

FIGURE 1. Proposed adaptive Adaptive-CS-Net architecture. The input and output of the network are stacks of three consequent knee MR images.

C. ADAPTIVE-CS-NETWORK

Starting from the network developed by Zhang and Ghanem, we developed the Adaptive-CS-Network approach. Our solution builds on top of the ISTA-Net solution based on three key innovations, here ordered by importance to the final network performance: i) the use of multi-scale and ii) multi-slice computations, together with iii) the introduction of soft MRI priors. We present them independently, building towards the update rule of the Adaptive-CS-Network model as presented in (16). Fig. 1 illustrates the proposed network.

First, many non-learned CS algorithms make use of **multi-scale** transforms to sparsify the signal. An example is given in (9), where $\mathbf{W}$ is a wavelet transform; a decomposition of the signal into a set of basis functions at different scales. We include this inductive bias in our design, and adopt a multi-scale transform $\mathcal{U}$, and its inverse $\hat{\mathcal{U}}$. As an additional design choice, we decide to sparsify and learn only the residual, therefore our update rule is written as follows:

$$\mathbf{x}_{i+1} = \hat{\mathcal{U}}(soft(\mathcal{U}(\mathbf{r}_{i+1}), \lambda_{s,f_s})) + \mathbf{r}_{i+1}, \tag{15}$$

where $\mathcal{U}$ comprises of 2D convolutions and non-linearities in the form of Leaky-ReLU to counteract the problem of dying neurons. To generate a multiscale representation, a max-pooling layer is used and the resulting features are then processed again by convolutional blocks and non-linearities. The exact design of $\mathcal{U}$ is presented in Fig. 1. The feature maps produced at the different scales are then thresholded using the soft-max function. Differently from ISTA-Net$^+$, we learn a lambda parameter and feature channel f_s for each scale s. This approach gives the network the flexibility of tuning the thresholds independently, hence reducing the complexity of the transforms learned by the convolutional operators. Finally, the filtered channels are transformed back into the image domain by the inverse $\hat{\mathcal{U}}$, consisting of

interpolation, 2D convolutions and Leaky-ReLU operators. Note that, contrary to the latest literature in deep learning networks, we decided not to adopt strided convolutions for sub- and up-sampling, which would increase the risk of creating checkerboard artifacts [39]; instead we took the more conservative approach of adopting pooling and interpolation layers for achieving better image quality. Overall, the computation represented by $\hat{\mathcal{U}}(soft(\mathcal{U}(\mathbf{r}_{i+1}), \lambda_{s,f_s}))$ is implemented with a UNet-like architecture [26], where the feature maps before the skip connections are filtered according to the parameter λ_{s,f_s}.

Second, it is important to note that the slice thickness of the dataset is much higher than the in-plane resolution. This indicates that inter-slice correlations are less useful for finer scales, and potentially damaging as they will become a confounder for the network. However, such information becomes beneficial at coarser scales, e.g., to facilitate the delineation of the bone in several slices. Since our transform $\mathcal{U}$ is multi-scale by nature, we found it beneficial to inject **neighboring slices** into the model, while leaving it to the network to identify at which scale the information will be used. To reduce the memory footprint of the model, we adopted a 2.5D convolution approach by concatenating neighbouring slices into the input tensor along the channel dimension, enabling to "reinvest" the saved GPU memory as compared to a truly 3D convolution approach, into more unrolled iterations. More details on the number of slices used and the definition of the loss function are given in Section IV-C.

Finally, we adopted a hybrid- or nudge- approach to incorporate additional **prior knowledge** into the reconstruction algorithm. We therefore computed additional information derived from the current estimate $\mathbf{x}_i$ together with k-space $M\mathbf{y}$. These soft priors, which are presented in the next section, capture some properties of an MR image that cannot be easily

learned by a deep neural network due to the limited size of the receptive field. The priors come in the form of images, and are provided as extra input channel to the transform $\mathcal{U}$. In this way, they are integrated in the computations performed by $\mathcal{U}$ whenever this is beneficial for the optimization of the loss function.

D. FINAL DESIGN

The overall update for a block B_{i+1} in the Adaptive-CS-Network model is defined as follows:

$$\begin{aligned} \mathbf{x}_{i+1} &= B_{i+1}(\mathbf{x}_i) \\ &= \mathbf{x}_i + \hat{\mathcal{U}}_i \left(soft \left(\mathcal{U}_i \left(\mathbf{x}_i, \mathbf{e}_{\mathrm{dc},i}, \mathbf{e}_{\phi,i}, \mathbf{e}_{\mathrm{bg},i} \right), \lambda_{s,f_s} \right) \right). \end{aligned} \quad (16)$$

Each block in the network learns different transforms $\mathcal{U}_i$ and $\hat{\mathcal{U}}_i$, enabling each block to focus on different properties of the image and effectively increasing the network capacity. Note that $\mathcal{U}_i$ and $\hat{\mathcal{U}}_i$ are different for every reconstruction block i.

In our final design, the transform $\mathcal{U}_i$ does not receive the data consistent image $\mathbf{r}_i$, as defined in (7), but rather the current estimate $\mathbf{x}_i$ together with the data consistency prior $\mathbf{e}_{\mathrm{dc},i}$ computed as follows:

$$\mathbf{e}_{\mathrm{dc},i} = \mathcal{F}^T (M\mathcal{F}\mathbf{x}_i - M\mathbf{y}). \quad (17)$$

This "soft data-consistency" update allows the network to evaluate the reliability of the acquired data and potentially compensate errors in the coil combination defined by $\mathcal{F}$ in (1).

The second prior we provided to the network, $\mathbf{e}_{\phi,i}$, represents the known phase response for spin-echo MR sequences. Theoretically, spin-echo sequences have zero phase everywhere in the image. In practice, however, slowly varying phase will occur, i.e. nonzero phase only in the low frequencies, due to hardware and acquisition imperfections. Taking this into account, it is noted that the final reconstructed image should be a real valued image after removal of the slowly varying phase. This information is captured in the following prior:

$$\mathbf{e}_{\phi,i} = \left\{ \mathbf{x}_i \cdot \frac{\mathbf{x}_{i,\mathrm{lpf}}^*}{\|\mathbf{x}_{i,\mathrm{lpf}}\|_2} \right\}_{imag}, \quad (18)$$

where $*$ denotes the complex conjugate, and lpf refers to low pass filtering. The low pass filter is chosen such that it corresponds to the center part of k-space which is fully sampled. By doing so, the low pass filtered image $\mathbf{x}_{i,\mathrm{lpf}}$ can be derived beforehand only once, hence $\mathbf{x}_{i,\mathrm{lpf}}$ is replaced by $\mathbf{x}_{0,\mathrm{lpf}}$.

Finally, we adopt a simple approach to estimate the location in x_i where the background is found, which is common in parallel imaging techniques. The following prior is applied:

$$\mathbf{e}_{\mathrm{bg},i} = \frac{\mathbf{x}_i}{\|\mathbf{x}_{i,\mathrm{lpf}}\|_2}. \quad (19)$$

This prior will penalize estimated signal content where $\|\mathbf{x}_{i,\mathrm{lpf}}\|$ is low, i.e., within the background. Again, $\mathbf{x}_{i,\mathrm{lpf}}$ is

replaced by $\mathbf{x}_{0,\mathrm{lpf}}$. Because $\mathbf{x}_{0,\mathrm{lpf}}$ is based on the fully measured central part of k-space, the image is artefact free albeit at low spatial resolution, leading to a reliable background identification.

In Fig. 1 the design of the Adaptive-CS-Network is shown, including the multi-scale transforms, the multi-slice computation and the priors provided as input. Note how the spin-echo and background priors are computed only for the central slice, in order to save GPU memory.

E. NETWORK TRAINING AND IMPLEMENTATION DETAILS

We implemented our models in PyTorch [40]. All the optimization experiments were performed on an NVIDIA V100 GPU with 16 GB RAM and the final network was trained on two NVIDIA V100 GPU with 16 GB RAM. In order to run as many experiments as possible given the challenge deadline, model optimization (see Section IV) was done with a relatively small model ($\leq$10 blocks), which we trained for 20 epochs. All the optimization networks were trained and validated on the highest acceleration rate of the challenge, i.e. 8$\times$ and for single-coil data, except for the number of the blocks which was performed for both 4$\times$ and 8$\times$, and for the priors which are more relevant for the multi-coil data. Since the ground truth for the test set was not available, all the quantitative comparisons were only done on the validation set.

For the challenge, we trained the final model using the training and validation datasets for 25 epochs and accelerations randomly selected from 2$\times$ to 10$\times$. The residual connections designed on a per-iteration basis, facilitates the learning and prevents the degradation of the error gradient throughout the architecture. The model was subsequently fine-tuned on eight data sub-populations identified by the acceleration (4$\times$ and 8$\times$), the protocol (PD and PDFS) and the scanner field strength (1.5T and 3T). Fine-tuning was then performed for 10 epochs on the sub-populations. This procedure was performed independently for the single- and multi-coil datasets, resulting in a total of 8 models. All models were trained using an exponentially decaying learning rate of 10^{-4}. The final models have 33M trainable parameters each; for the single-coil data this leads to an inference time of approximately 327 ms, while it takes approximately 518 ms to compute the reconstruction of a multi-coil dataset on an NVIDIA V100 GPU.

IV. EXPERIMENTS AND RESULTS: MODEL OPTIMIZATION

In this section we present how we optimized the network configuration, on a smaller model with 10 reconstruction blocks, using the quantitative measures reported in Section II-B for validation. We performed experiments on the number of the blocks, the loss functions, the influence of using adjacent slices, the optimizer, and the soft priors. A repeated measure one-way ANOVA test was performed on the SSIM values using a significance level of $p = 0.05$. P-values are only stated for the comparisons between the best method and

FIGURE 2. The effect of the number of blocks on performance, using the 4× and 8× single-coil validation data. The variance values are shown by the bars. The stars in the first plot show one-way ANOVA statistical significance.

the other methods. In all the experiments a learning rate of 0.0001 was used.

A. NUMBER OF BLOCKS

The proposed model consists of multiple blocks, related to the number of unrolled iterations of the ISTA scheme. Increasing the number of blocks leads to an increase in the number of parameters of the model, and subsequently training time and GPU memory usage as well as an increase in risk of overfitting. In this experiment we investigated the effect of the number of the blocks on the quality of reconstructed images. Tests were ran with the 2D network for 4× and 8× acceleration rates without neighboring slices, MSE as loss function, RMSprop as optimizer, and with the Unet-like architecture of 16 filter maps for each convolutional layer. Fig. 2 reports the relative changes to a single block of our quantitative metrics. Based on the experiments, increasing the number of the blocks will improve the performance of the network. Therefore, the final network was configured with the maximum number of blocks that could be fitted into GPU memory: 25 blocks. However, for the optimization experiments below only 10 blocks were employed to limit the duration of the training.

B. LOSS FUNCTIONS

In this experiment we investigated the effect of a wide range of differentiable loss functions on the performance of our network. Here, we used the single slice reconstruction network with only 10 blocks, RMSprop as the optimizer, and 16 filter maps for each convolutional layer. The models were trained for 20 epochs to ensure convergence of the model. The evaluated loss functions included MSE, perceptual loss (PL) [41], ℓ_1, Huber [42] and multi-scale structural similarity index (MSSIM) [43]. The PL loss function was calculated using a pre-trained VGG-16 at layers relu1_2, relu2_2, and relu3_3.

MSSIM [43] builds upon SSIM (see Section II-B3) by incorporating structural similarity at multiple image resolutions, thereby supplying more flexibility compared to SSIM, and is defined as follows:

$$\text{MSSIM} = \left[l_M(\mathbf{r^c}, \mathbf{t^c})\right]^{\alpha_M} \prod_{i=1}^{M} \left[c_i(\mathbf{r^c}, \mathbf{t^c})\right]^{\beta_i} \left[s_i(\mathbf{r^c}, \mathbf{t^c})\right]^{\gamma_i}, \quad (20)$$

TABLE 1. The effect of the loss function on performance, using the 8× single-coil validation data. Stars denote one-way ANOVA statistical significance.

Loss function	SSIM $\mu \pm \sigma$	NMSE $\mu \pm \sigma$	PSNR $\mu \pm \sigma$	p-value
MSE	0.657±0.149*	0.046±0.029	30.2±2.8	≪0.001
Perceptual loss	0.664±0.157*	0.061±0.044	29.2±3.2	≪0.001
Huber	0.664±0.148*	0.062±0.041	29.1±3.0	≪0.001
ℓ_1	0.664±0.148*	0.062±0.041	29.1±3.0	≪0.001
SSIM	0.662±0.145*	0.065±0.041	28.9±2.8	≪0.001
MSSIM [43]	0.671±0.143*	0.050±0.034	30.1±3.1	≪0.001
Eq. (21)	**0.673±0.143**	**0.048±0.033**	**30.3±3.1**	

where $\mathbf{r}^c$, $\mathbf{t}^c$, denote the reconstructed and target images respectively, M is the number of scales used, l_M, c_i and s_i are the luminance, contrast, and structure as defined in [36], α_M, β_i, and γ_i are the weights of the distortion factors at different resolution levels. We adopted the same weights as reported in [43].

Zhao *et al.* [44] reported that a linear combination of SSIM and ℓ_1 preserves the different properties of an image better than each separately: SSIM encourages the network to preserve structural and contrast information, while ℓ_1 enforces sensitivity to uniform biases to preserve luminance [45]. Since MSSIM reached higher metric values than SSIM (see Table 1), we deployed a weighted summation of MSSIM [43] and ℓ_1:

$$L = \alpha\text{MSSIM}(\mathbf{r^c}, \mathbf{t^c}) + (1 - \alpha)\|\mathbf{r^c} - \mathbf{t^c}\|_1, \quad (21)$$

where $\alpha = 0.84$ was chosen, following Zhao *et al.* [44]. Note that, compared to the ISTA-Net approach, we found it beneficial not to adopt the discrepancy loss as presented in Eq. (12) for two reasons. First, we empirically found that tuning the loss multiplier θ is not straightforward, leading to sub-optimal results in terms of the reconstruction loss. Secondly, computing the discrepancy loss is very demanding in terms of GPU memory, requiring to perform a second forward pass where only the thresholding operation is ignored. While feasible, it requires to make the model significantly smaller in terms of learnable parameters, hence reducing model performance significantly.

Table 1 reports the quantitative results for the different loss functions. The weighted linear combination of MSSIM and ℓ_1 yielded the best results, where the p-values indicate that the improvement achieved thanks to our modifications is highly consistent across all scans, despite the small improvements on SSIM-values. Fig. 3 shows two example results for the different loss functions, confirming the favorable results for the model trained using a combination of MSSIM and ℓ_1. Therefore, this loss function was selected for training the final model. For the remainder of the experiments, MSSIM is used as loss function.

C. MULTI-SLICE NETWORK

The resolution of the images in the dataset is anisotropic with a voxel size of $0.5 \times 0.5 \times 3 \ mm^3$. Due to the correlation between adjacent slices with respect to anatomical structures in MRI images, we performed an experiment to assess

Target	ℓ_1	SSIM	MSSIM	MSSIM + ℓ_1

FIGURE 3. Two examples of single-coil 4× for the different loss functions. A small network is used to test several losses. SSIM values are shown in yellow.

TABLE 2. The effect of adopting a 2.5D approach on the 8× single-coil data using the small model. W denotes the loss weight applied to the neighboring slices. Stars denote one-way ANOVA statistical significance.

Network	SSIM $\mu \pm \sigma$	NMSE $\mu \pm \sigma$	PSNR $\mu \pm \sigma$	*p*-value
2D	0.671±0.143*	0.050±0.034	30.1±3.1	≪0.001
2.5D W0.1	0.549±0.128*	0.089±0.034	26.8±2.1	≪0.001
2.5D W0.2	0.548±0.128*	0.090±0.033	26.8±2.1	≪0.001
2.5D	**0.674±0.143**	**0.048±0.033**	**30.3±3.1**	

TABLE 3. The effect of the optimizer on performance, using the 8× single-coil validation data. Stars denote one-way ANOVA statistical significance.

Optimizer	SSIM $\mu \pm \sigma$	NMSE $\mu \pm \sigma$	PSNR $\mu \pm \sigma$	*p*-value
RMSprop	0.673±0.143*	0.048±0.033	30.3±3.1	≪0.001
LookAhead	0.668±0.140*	0.050±0.032	30.0±2.9	≪0.001
Ranger	0.668±0.140*	0.050±0.032	30.0±2.9	≪0.001
RAdam	**0.674±0.141**	**0.048±0.032**	**30.3±3.0**	

whether inclusion of neighbouring slices into the reconstruction might improve the performance. We compared the 2D scheme using only the center slice with three alternative 2.5D schemes: i) the neighboring slices were used together with the center slice as input, but only the center slice was used in the loss function (network 2.5D); ii) and iii) the neighboring slices are also used in the loss, with different weights (0.1 vs 0.2 for the neighbors; 1.0 for the center slice). To compute the first and last slice, we pad the volume with replicas of the edge slices. MSSIM was used for the loss function, 10 blocks, RMSprop as the optimizer, and 16 feature maps.

Table 2 shows the results of this experiment, showing that the 2.5D schema very consistently improves over the 2D scheme, and that the loss should only be defined on the center slice. For the final model, this scheme was selected.

D. OPTIMIZER

We experimented with different optimizers including RMSprop, rectified Adam (RAdam) [46], LookAhead [47] and Ranger [48]. RAdam exploits a dynamic rectifier to adjust the adaptive momentum of Adam [49]. LookAhead not only uses an adaptive learning rate and accelerated schemes but also iteratively updates two sets of weights, i.e. fast and slow weights. Ranger combines Radam and LookAhead optimizers into a single one. We used the 2D network with 10 blocks and 16 feature maps for each layer, and MSSIM the loss function.

Table 3 tabulates the results for the different optimizers. Since the best results were obtained for the RAdam optimizer, very consistently improving over the other optimizers, this was used for the final network.

E. ADAPTIVE-CS-NET VS ISTA-Net$^+$

In this experiment, we compare the proposed model to ISTA-Net$^+$ [33]. For this experiment, a 2D network with 10 blocks and 16 feature maps per layer was used, SSIM as loss function, and RAdam as the optimizer. Since ISTA-Net$^+$ uses a much smaller single scale architecture with much fewer network parameters, we added an experiment increasing the feature maps for ISTA-Net$^+$ such that the number of parameters was the same as for our architecture. According to the results reported in Table 4, the proposed model outperforms ISTA-Net$^+$ significantly. Figure 4 shows a qualitative comparison between ISTA-Net$^+$ and Adaptive-CS-Net on the single-coil 4× dataset. Although for the first image Adaptive-CS-Net reconstructed a better output in terms of the anatomical structure, the output of ISTA-Net-L$^+$ has a higher SSIM value. This implies that the radiological evaluation is a complementary step to judge the quality of the results.

F. SOFT PRIORS

To assess the contribution of the additional soft priors, we compared the full model against a version without known phase behaviour $\mathbf{e}_{\phi,i}$ and without background

FIGURE 4. Qualitative comparison of Adaptive-CS-Net vs ISTA-Net$^+$ on the single-coil 4$\times$ dataset. The SSIM values are shown in yellow.

TABLE 4. Adaptive-CS-Net vs ISTA-Net$^+$ on the 8$\times$ single-coil dataset. ISTA-Net$^+$ has 0.75M trainable parameters, while ISTA-Net-L$^+$ and A-CS-Net have 2.12M trainable parameters. Stars denote one-way ANOVA statistical significance.

Model	SSIM $\mu \pm \sigma$	NMSE $\mu \pm \sigma$	PSNR $\mu \pm \sigma$	p-value
ISTA-Net$^+$	0.547±0.117*	0.169±0.022	23.8±1.9	≪0.001
ISTA-Net-L$^+$	0.543±0.119*	0.103±0.038	26.2±2.0	≪0.001
A-CS-Net	**0.671±0.143**	**0.050±0.034**	**30.1±3.1**	

TABLE 5. The effect of adding priors to the final network on performance, using the multi-coil test data.

Acceleration	prior	SSIM	NMSE	PSNR
4x	−	0.772	**0.025**	30.98
	+	**0.773**	0.028	**33.49**
8x	−	0.674	**0.038**	**30.90**
	+	**0.675**	0.044	30.27

information $\mathbf{e}_{\text{bg},i}$. Visually, we observed only small differences. To verify the differences in a realistic setting, we submitted the results to the public leaderboard of the fastMRI challenge. As shown in Table 5, the network with all priors performed better in terms of the SSIM metric, although the results worsened in terms of NMSE and PSNR. Despite the fact that the improvement was minimal, we decided to adopt all priors for the final model to ensure our participation in the last challenge phase, since the selection was based on SSIM.

V. ADAPTIVE-CS-NET: SUBMITTED MODEL

In this section, we describe the configuration of the submitted model [50] and analyze the resulting reconstructions. The final performance is evaluated with the quantitative metrics on the test and challenge datasets, and by presenting the radiological scores for the challenge dataset as performed by the fastMRI challenge organizers.

Following our model optimization study, the configuration of the final model was determined as follows. The linear combination of MSSIM and ℓ_1 (21) was chosen as the loss function. The 2.5D scheme was chosen with two neighboring slices, with the loss applied only on the central slice. For training the model, the RAdam optimizer was deployed. Fig. 5 shows the structure of the final network. Each block is determined by three parameters for the denoiser: 1) the number of scales for the denoiser $\mathcal{U}, \hat{\mathcal{U}}$, 2) the kernel size used in the convolutions and, 3) the number of feature maps in the first convolutional layer, which is then doubled at each scale. According to the experiments presented in Fig. 2, the number of reconstruction blocks greatly affects the reconstruction performance, empirically observing that performance still improves when 15 blocks are used. The available GPU memory is a limiting factor when designing a deep neural network. To allow for a large number of blocks, we chose a different design in each block, mixing a less powerful design (16 filters) with more powerful ones (64 filters). By adopting this strategy, our final design contained 25 reconstruction blocks and has 33M parameters.

Fig. 6 shows example results of the final network for the multi-coil track from the validation dataset. Fig. 7 shows examples from the test and challenge datasets. Table 6 shows

FIGURE 5. Final network, each block has the same structure as shown in Fig. 1 and is defined by $\mathcal{U}, \hat{\mathcal{U}}$[number of scales, kernel size, number of feature maps in the first scale]. For all layers Leaky ReLU was used as the activation function.

FIGURE 6. Example results of the final model for the multi-coil track accelerated by 8× on the validation dataset. Top row depicts the target image, bottom row the reconstructed images with the SSIM value in yellow.

the SSIM, NMSE, and PSNR values for the test and challenge set (as described in Section II-A), for the images with and without fat suppression and both combined, for both single- and multi-coil MRI scans. For the radiological evaluation, our method scored 2.285, 1.286, and 2.714 for multi-coil 4×, multi-coil 8×, and single-coil 4×, respectively (the closer to 1, the better). The average runtimes for the model are 518 and 327 milliseconds for the multi-coil and the single-coil data, respectively. More details on the results for the challenge were presented in [34].

VI. DISCUSSION

In this paper we propose a general method, named Adaptive-CS-Net, for reconstructing undersampled MRI data, combining ideas from compressed sensing theory with ideas from MR physics and deep learning. The method was developed in the context of the 2019 fastMRI challenge, which focused on accelerating knee MR imaging. The proposed network is an unrolled iterative learning-based reconstruction scheme, in which a large number of reconstruction blocks refine the MR image by denoising the signal in a learned and multi-scale fashion. Moreover, we added neighboring slices as input to the sparsifying transform, as well as a number of soft priors that encode MRI domain knowledge.

The main driver of the performance of our network is the multi-scale architecture, as demonstrated in a direct comparison with ISTA-Net$^+$ that is corrected by the number of trainable parameters. According to the experimental results

TABLE 6. Results for the final model for single- and multi-coil data on the test and challenge dataset.

Dataset	Coil		Detail	SSIM	NMSE	PSNR
Test	multi	4x	ALL	0.928	0.005	39.9
			PD	0.961	0.002	41.7
			PDFS	0.891	0.009	37.9
		8x	ALL	0.888	0.009	36.8
			PD	0.937	0.005	38.5
			PDFS	0.843	0.013	35.3
	single	4x	ALL	0.777	0.027	33.7
			PD	0.877	0.010	36.9
			PDFS	0.685	0.043	30.7
		8x	ALL	0.680	0.042	30.5
			PD	0.777	0.019	32.4
			PDFS	0.575	0.067	28.5
Challenge	multi	4x	ALL	0.927	0.005	39.9
		8x	ALL	0.901	0.009	37.4
	single	4x	ALL	0.751	0.030	32.7

on the number of blocks for 4× and 8× accelerations of both single- and multi-coil data, we showed that the number of blocks has a large impact on model performance. Therefore, it was decided to use the maximum number of blocks that we could fit into the GPU memory, where we adopted different model designs for the different blocks to save memory. It might be expected that beyond a certain number of blocks, overfitting of the data might occur. However, signs of overfitting were not observed during training and the final number of blocks was only marginally larger than tested in the optimization experiments. Whether further increase in the number of blocks could result in even better

FIGURE 7. Example results of the final network from the test and challenge datasets, for which no ground truth reconstructions are available.

performance could be the topic of further experiments. This would, however, need better hardware, as the current design is memory- and time-bound during training. With the current configuration, final model training took approximately 7 days on two V100 GPUs.

We experimented with a large variety of loss functions. Results showed that the linear summation of MSSIM and ℓ_1 performed best. Figure 6 showed that poor SNR data yield very low SSIM scores. Surprisingly, within high SNR data, a large variance of SSIM scores is also found. This highlights the fact that further research is required in order to develop better quality metrics. Moreover, we defined a 2.5D scheme to train the network in which three adjacent slices were reconstructed while the loss function was calculated only for the central slice. The proposed scheme outperformed the 2D network as well as 2.5D networks in which the loss was calculated over all slices. By incorporating the neighbouring slices, the network can exploit existing correlations into the reconstruction of the target slice, which is our main target as defined by our loss. It can be expected that for MRI acquisition with less asymmetric voxel sizes, the inclusion of information of neighbouring slices would become more important. However, weighing the loss of the neighbouring slices resulted in less optimal results since it forces the network to solve a more difficult problem: the network has to reconstruct multiple slices instead of a single one. This reduces the effective network capacity per slice, leading to a degradation of the reconstruction performance. We tested different optimizers, where the newly introduced RAdam outperformed the others and we used it for training the final network. We also incorporated prior knowledge, including data consistency, known phase behaviour and background discrimination to support the network in the reconstruction process. We observed that these priors provided only limited extra performance to the network, resulting in visually similar images and minimal difference in the metrics.

We can conclude that the Adaptive-CS-Net is sufficiently powerful to learn directly from the data how to reconstruct the undersampled k-space, being the multi-scale structure and the use of many reconstruction blocks the main driver of our performance. As a future work, we want to better understand how much the network is relying on the priors by adopting interpretable AI techniques such as differentiable image parameterizations for feature visualization [51]. Stronger use of the priors via the loss function is an additional option.

As mentioned before, the radiologist scores were based on the visual quality of the reconstructed images and not on diagnostic interchangeability. Therefore, designing a network based on the diagnosis can be considered a point for further research. We furthermore observed that optimizing for SSIM was needed for reaching the final stage of the challenge, but is not necessarily an ideal representative of radiological image quality. This observation was very recently confirmed in a comparative study by others [52]. The proposed method outperforms the benchmark networks, including U-net [26], cascade net [27], KIKI-net [28], and PD-net [29], on the single-coil track as reported in [24]. It outperforms as well the i-RIM model [16] on the Multi-coil track but not the single coil track [34].

VII. CONCLUSION

In this paper we propose an adaptive intelligence algorithm called Adaptive-CS-Net, which was developed in the context of the 2019 fastMRI challenge. In the two clinically relevant tracks of the challenge, using multi-coil MRI acquisitions, the proposed method was leading, while on a simulated single-coil track the method ranked 3rd.

Jeroen Hendrikus Fransiscus van Gemert contributed equally to this work.) (Elwin de Weerdt and Marius Staring are co-senior authors.)

REFERENCES

[1] M. Zaitsev, J. Maclaren, and M. Herbst, "Motion artifacts in MRI: A complex problem with many partial solutions," *J. Magn. Reson. Imag.*, vol. 42, no. 4, pp. 887–901, Oct. 2015.

[2] M. S. Cohen and R. M. Weisskoff, "Ultra-fast imaging," *Magn. Reson. Imag.*, vol. 9, no. 1, pp. 1–37, Jan. 1991.

[3] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[4] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, 2007.

[5] E. J. Candès, Y. C. Eldar, D. Needell, and P. Randall, "Compressed sensing with coherent and redundant dictionaries," *Appl. Comput. Harmon. Anal.*, vol. 31, no. 1, pp. 59–73, Jul. 2011.

[6] Z.-P. Liang, F. Boada, R. T. Constable, E. M. Haacke, P. C. Lauterbur, and M. R. Smith, "Constrained reconstruction methods in MR imaging," *Rev. Magn. Reson. Med.*, vol. 4, no. 2, pp. 67–185, 1992.

[7] K. P. Pruessmann, M. Weiger, M. B. Scheidegger, and P. Boesiger, "SENSE: Sensitivity encoding for fast MRI," *Magn. Reson. Med.*, vol. 42, no. 5, pp. 952–962, Nov. 1999.

[8] D. Liang, J. Cheng, Z. Ke, and L. Ying, "Deep MRI reconstruction: Unrolled optimization algorithms meet neural networks," 2019, *arXiv:1907.11711*. [Online]. Available: http://arxiv.org/abs/1907.11711

[9] T. M. Quan, T. Nguyen-Duc, and W.-K. Jeong, "Compressed sensing MRI reconstruction using a generative adversarial network with a cyclic loss," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1488–1497, Jun. 2018.

[10] 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 Trans. Med. Imag.*, vol. 38, no. 1, pp. 167–179, Jan. 2019.

[11] Y. Guo, C. Wang, H. Zhang, and G. Yang, "Deep attentive wasserstein generative adversarial networks for MRI reconstruction with recurrent context-awareness," 2020, *arXiv:2006.12915*. [Online]. Available: http://arxiv.org/abs/2006.12915

[12] J. Schlemper, G. Yang, P. Ferreira, and A. Scott, "Stochastic deep compressive sensing for the reconstruction of diffusion tensor cardiac MRI," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, 2018, pp. 295–303.

[13] 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 Trans. Med. Imag.*, vol. 37, no. 6, pp. 1310–1321, Jun. 2018.

[14] S. Yu, H. Dong, G. Yang, G. Slabaugh, P. Luigi Dragotti, X. Ye, F. Liu, S. Arridge, J. Keegan, D. Firmin, and Y. Guo, "Deep de-aliasing for fast compressive sensing MRI," 2017, *arXiv:1705.07137*. [Online]. Available: http://arxiv.org/abs/1705.07137

[15] P. Putzky, D. Karkalousos, J. Teuwen, N. Miriakov, B. Bakker, M. Caan, and M. Welling, "I-RIM applied to the fastMRI challenge," 2019, *arXiv:1910.08952*. [Online]. Available: http://arxiv.org/abs/1910.08952

[16] P. Putzky and M. Welling, "Invert to learn to invert," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 446–456.

[17] 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, Mar. 2018.

[18] D. Lee, J. Yoo, S. Tak, and J. C. Ye, "Deep residual learning for accelerated MRI using magnitude and phase networks," *IEEE Trans. Biomed. Eng.*, vol. 65, no. 9, pp. 1985–1995, Sep. 2018.

[19] G. Wang, J. C. Ye, K. Mueller, and J. A. Fessler, "Image reconstruction is a new frontier of machine learning," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1289–1296, Jun. 2018.

[20] 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," *Magn. Reson. Med.*, vol. 79, no. 6, pp. 3055–3071, Jun. 2018.

[21] J. Sun, "Deep ADMM-Net for compressive sensing MRI," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 10–18.

[22] J. Yang, Y. Zhang, and W. Yin, "A fast alternating direction method for TVL1-L2 signal reconstruction from partial Fourier data," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 288–297, Apr. 2010.

[23] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-based deep learning architecture for inverse problems," *IEEE Trans. Med. Imag.*, vol. 38, no. 2, pp. 394–405, Feb. 2019.

[24] Z. Ramzi, P. Ciuciu, and J.-L. Starck, "Benchmarking deep nets MRI reconstruction models on the fastmri publicly available dataset," in *Proc. IEEE 17th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2020, pp. 1441–1445.

[25] J. Zbontar et al., "FastMRI: An open dataset and benchmarks for accelerated MRI," 2018, *arXiv:1811.08839*. [Online]. Available: https://arxiv.org/abs/1811.08839

[26] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, 2015, pp. 234–241.

[27] S. Ramanarayanan, B. Murugesan, K. Ram, and M. Sivaprakasam, "DC-WCNN: A deep cascade of wavelet based convolutional neural networks for MR image reconstruction," in *Proc. IEEE 17th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2020, pp. 647–658.

[28] 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," *Magn. Reson. Med.*, vol. 80, no. 5, pp. 2188–2201, Nov. 2018.

[29] J. Adler and O. Oktem, "Learned primal-dual reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1322–1332, Jun. 2018.

[30] J. Caballero, A. N. Price, D. Rueckert, and J. V. Hajnal, "Dictionary learning and time sparsity for dynamic MR data reconstruction," *IEEE Trans. Med. Imag.*, vol. 33, no. 4, pp. 979–994, Apr. 2014.

[31] A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *J. Math. Imag. Vis.*, vol. 40, no. 1, pp. 120–145, May 2011.

[32] M. Seitzer, "Adversarial and perceptual refinement for compressed sensing mri reconstruction," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, 2018, pp. 232–240.

[33] J. Zhang and B. Ghanem, "ISTA-net: Interpretable optimization-inspired deep network for image compressive sensing," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1828–1837.

[34] 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," *Magn. Reson. Med.*, vol. 84, no. 6, pp. 3054–3070, Dec. 2020.

[35] F. Knoll, "FastMRI: A publicly available raw k-Space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning," *Radiol., Artif. Intell.*, vol. 2, no. 1, Jan. 2020, Art. no. e190007.

[36] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.

[37] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, Jan. 2009.

[38] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proc. Nat. Acad. Sci. USA*, vol. 106, no. 45, pp. 18914–18919, Nov. 2009.

[39] A. Araujo, W. Norris, and J. Sim, "Computing receptive fields of convolutional neural networks," *Distill*, vol. 4, no. 11, p. 151, Nov. 2019.

[40] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," in *Proc. NIPS*, 2017, pp. 1–9.

[41] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 694–711.

[42] P. J. Huber, "Robust estimation of a location parameter," in *Breakthroughs in Statistics*. New York, NY, USA: Springer, 1992, pp. 492–518.

[43] Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multiscale structural similarity for image quality assessment," in *Proc. 37th Asilomar Conf. Signals, Syst. Comput.*, vol. 2, Nov. 2003, pp. 1398–1402.

[44] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, "Loss functions for image restoration with neural networks," *IEEE Trans. Comput. Imag.*, vol. 3, no. 1, pp. 47–57, Mar. 2017.

[45] H. Zhao, O. Gallo, I. Frosio, and J. Kautz, "Loss functions for neural networks for image processing," 2015, *arXiv:1511.08861*. [Online]. Available: http://arxiv.org/abs/1511.08861

[46] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the variance of the adaptive learning rate and beyond," 2019, *arXiv:1908.03265*. [Online]. Available: http://arxiv.org/abs/1908.03265

[47] M. Zhang, J. Lucas, J. Ba, and G. E. Hinton, "Lookahead optimizer: K steps forward, 1 step back," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 9593–9604.

[48] Q. Tong, G. Liang, and J. Bi, "Calibrating the adaptive learning rate to improve convergence of ADAM," 2019, *arXiv:1908.00700*. [Online]. Available: http://arxiv.org/abs/1908.00700

[49] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent.*, 2015, pp. 1–9.

[50] N. Pezzotti, E. de Weerdt, S. Yousefi, M. S. Elmahdy, J. van Gemert, C. Schälke, M. Doneva, T. Nielsen, S. Kastryulin, B. P. F. Lelieveldt, M. J. P. van Osch, and M. Staring, "Adaptive-CS-net: FastMRI with adaptive intelligence," 2019, *arXiv:1912.12259*. [Online]. Available: http://arxiv.org/abs/1912.12259

[51] A. Mordvintsev, N. Pezzotti, L. Schubert, and C. Olah, "Differentiable image parameterizations," *Distill*, vol. 3, no. 7, p. e12, Jul. 2018.

[52] 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 Trans. Med. Imag.*, vol. 39, no. 4, pp. 1064–1072, Apr. 2020.

NICOLA PEZZOTTI received the B.Sc. and M.Sc. degrees in computer science and engineering from the University of Brescia, Italy, in 2009 and 2011, respectively, and the Ph.D. degree *(cum laude)* from the Delft University of Technology, The Netherlands, in 2018. He was a Research and Development Engineer with Open Technology S.r.l. from 2011 to 2014, and a Visiting Scientist with INRIA Saclay, Paris, in 2017, and Google AI, Zürich, in 2018. He is currently a Senior Scientist with Philips Research, Eindhoven, The Netherlands, and an Assistant Professor with the Eindhoven University of Technology. His research interests include machine learning, visual analytics, explainable AI, optimization techniques, and software engineering. He was a recipient of several awards, including the IEEE VGTC Best Dissertation Award, the TU Delft Excellence in Research, and the Dirk Bartz Prize for Visual Computing in Medicine.

SAHAR YOUSEFI received the B.S. degree in software engineering from Alzahra University, Tehran, Iran, in 2008, the M.S. degree in artificial intelligence from the Shahrood University of Technology, Shahrood, Iran, in 2009, and the Ph.D. degree in artificial intelligence from the Sharif University of Technology, Tehran, in 2018. From 2017 to 2018, she was a Research Visitor with the Division of Image Processing, Leiden University Medical Center (LUMC), Leiden, The Netherlands. In 2018, she started her career in the field of image processing using deep learning at LUMC as a deep learning Researcher. Her research interests include deep learning, machine learning, and image and video processing.

MOHAMED S. ELMAHDY received the B.S. and M.S. degrees in biomedical engineering from Cairo University, Egypt, in 2013 and 2017, respectively. He is currently pursuing the Ph.D. degree in biomedical engineering with the Leiden University Medical Center, Leiden, The Netherlands. His M.S. thesis focused on subvocal speech recognition using deep learning. From 2013 to 2017, he was a Teaching and Research Assistant with the Faculty of Engineering, Cairo University. His research interests include developing medical image registration and segmentation algorithms using deep learning, image reconstruction, and multi task learning.

JEROEN HENDRIKUS FRANSISCUS VAN GEMERT was born in Nijmegen, The Netherlands, in 1986. He received the bachelor's and master's degrees *(cum laude)* in electrical engineering from the Delft University of Technology, in 2011 and 2013, respectively. During his M.Sc. research, he worked with the company Lismar Engineering B.V. on crack detection for plate rolls used for pressing steel. Subsequently, he started to work as a Software Engineer for the company Imtech Traffic & Infra, where he worked on the power management systems for the trams in The Hague and Amsterdam. After a year, he returned to Delft for his Ph.D. led by an Associate Professor Rob Remis at the Circuits & Systems Group, EEMCS Faculty, headed by Prof. Alle-Jan van der Veen. In 2015, he started his Ph.D. research and closely collaborated with the C.J. Gorter Center for High-field MRI, Leiden University Medical Center, headed by Prof. A. Webb, and the Department of Microwave Sensing, Signals and Systems, TU Delft, headed by Prof. A. Yarovoy. He is currently working with Philips Research as a MRI Reconstruction Engineer.

CHRISTOPHE SCHUELKE graduated from the Ecole Polytechnique, France, in 2011. He received the M.S. degree in physics from ETH Zürich, Switzerland, in 2012, and the Ph.D. degree in physics from the UniversitéParis 7 and University of Roma La Sapienza, in 2016. Since 2016, he has been a Research Scientist with Philips Research, Hamburg, Germany. His research interests include MRI image reconstruction, signal processing, information theory, and machine learning.

MARIYA DONEVA (Member, IEEE) received the B.Sc. and M.Sc. degrees in engineering physics from the University of Oldenburg, Germany, in 2006 and 2007, respectively, and the Ph.D. degree from the Institute for Signal Processing, University of Luebeck, in 2010. She was a Research Associate with the Department of Electrical Engineering and Computer Sciences, UC Berkeley, from 2015 to 2016. She is currently a Senior Scientist with the Philips Research, Hamburg, Germany. Her research interests include methods for efficient data acquisition, image reconstruction, and quantitative parameter mapping in the context of magnetic resonance imaging. She was a recipient of the Junior Fellow Award of the International Society for Magnetic Resonance in Medicine in 2011.

TIM NIELSEN received the Ph.D. degree in physics from the University of Bielefeld. He is currently a Senior Scientist with the Philips Research, Hamburg, Germany. His research interests include methods for efficient data acquisition, image reconstruction, and machine learning in the context of magnetic resonance imaging.

SERGEY KASTRYULIN was born in Rostov-on-Don, Russia, in 1993. He received the specialist degree in organic chemistry and the M.S. degree in computer science from Southern Federal University, Rostov-on-Don, in 2015 and 2019, respectively. He is currently pursuing the Ph.D. degree in artificial intelligence with the Skolkovo Institute of Science and Technology, Moscow, Russia. From 2018 to 2019, he was a Data Scientist with uKit.group, working on generative adversarial networks for image generation. Since 2019, he has been a Research Scientist with Philips Research, Russia. His research interests include investigating inverse problems in medical imaging using generative modeling, compression of artificial neural networks, and model understanding in the field of artificial neural networks.

BOUDEWIJN P. F. LELIEVELDT (Member, IEEE) received the Ph.D. degree in medical image analysis from Leiden University, in 1999. He is currently heading the Division of Image Processing, Leiden University Medical Center, and hold a Medical Delta Professor Chair of biomedical imaging with Leiden University and the Delft University of Technology. His research interests include dimensionality reduction methods, with application in complex biomedical datasets.

MATTHIAS J. P. VAN OSCH received the M.Sc. degree in applied physics from the University of Twente, Enschede, The Netherlands, in 1997, and the Ph.D. degree from the Image Sciences Institute, University Medical Center, Utrecht, The Netherlands, in 2002. From 2002 to 2004, he was a Postdoctoral Researcher with the Image Sciences Institute, University Medical Center, Utrecht. In 2004, he moved to the Leiden University Medical Center, Leiden, The Netherlands, at which he became a Full Professor on Experimental Cerebrovascular Imaging in 2019. His research interests include the development of new MRI sequences with a special focus on hemodynamic measurements and clinical brain applications of ultra-high field MRI. He is currently an Active Member of the ESMRMB and ISMRM.

ELWIN DE WEERDT received the B.Sc. and M.Sc. degrees in aerospace engineering and the Ph.D. degree from the Delft University of Technology, The Netherlands, in 2003, 2005, and 2010, respectively. He was a Researcher with the Aerospace Software and Technology Institute (ASTI) from 2006 to 2011, and a Research Intern with the National Aerospace Laboratory (NLR) from 2003 to 2004. He is currently a Reconstruction Algorithm Architect with Philips Healthcare, Best, The Netherlands. His research interests include methods for efficient and robust data acquisition and image reconstruction in the context of magnetic resonance imaging.

MARIUS STARING received the M.Sc. degree in applied mathematics from the University of Twente, in 2002, and the Ph.D. degree from the UMC Utrecht, in 2008. Since then, he has been with LUMC, and since 2015, he has also been a part-time with TU Delft. He is currently an Associate Professor of medical image analysis, working with the Leiden University Medical Center, where he leads the Biomedical Machine Learning research line. His research interests include image registration and all image analysis aspects around radiotherapy, image acquisition and radiology. He is also an Associate Editor of IEEE Transactions on Medical Imaging (TMI), and serves as a program committee member of several international conferences.

• • •