Parallel-stream fusion of scan-specific and scan-general priors for learning deep MRI reconstruction in low-data regimes

Salman Ul Hassan Dar [a,i,1], Şaban Öztürk [b,c,1], Muzaffer Özbey [d], Kader Karli Oguz [e,f], Tolga Çukur [b,f,g,h,*]

[a] Department of Internal Medicine III, Heidelberg University Hospital, 69120, Heidelberg, Germany
[b] Department of Electrical and Electronics Engineering, Bilkent University, Ankara 06800, Turkey
[c] Department of Electrical-Electronics Engineering, Amasya University, Amasya 05100, Turkey
[d] Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign, IL 61820, United States
[e] Department of Radiology, University of California, Davis, CA 95616, United States
[f] Department of Radiology, Hacettepe University, Ankara, Turkey
[g] National Magnetic Resonance Research Center (UMRAM), Bilkent University, Ankara 06800, Turkey
[h] Neuroscience Graduate Program, Bilkent University, Ankara 06800, Turkey
[i] AI Health Innovation Cluster, Heidelberg, Germany

ARTICLE INFO

ABSTRACT

Magnetic resonance imaging (MRI) is an essential diagnostic tool that suffers from prolonged scan times. Reconstruction methods can alleviate this limitation by recovering clinically usable images from accelerated acquisitions. In particular, learning-based methods promise performance leaps by employing deep neural networks as data-driven priors. A powerful approach uses scan-specific (SS) priors that leverage information regarding the underlying physical signal model for reconstruction. SS priors are learned on each individual test scan without the need for a training dataset, albeit they suffer from computationally burdening inference with nonlinear networks. An alternative approach uses scan-general (SG) priors that instead leverage information regarding the latent features of MRI images for reconstruction. SG priors are frozen at test time for efficiency, albeit they require learning from a large training dataset. Here, we introduce a novel parallel-stream fusion model (PSFNet) that synergistically fuses SS and SG priors for performant MRI reconstruction in low-data regimes, while maintaining competitive inference times to SG methods. PSFNet implements its SG prior based on a nonlinear network, yet it forms its SS prior based on a linear network to maintain efficiency. A pervasive framework for combining multiple priors in MRI reconstruction is algorithmic unrolling that uses serially alternated projections, causing error propagation under low-data regimes. To alleviate error propagation, PSFNet combines its SS and SG priors via a novel parallel-stream architecture with learnable fusion parameters. Demonstrations are performed on multi-coil brain MRI for varying amounts of training data. PSFNet outperforms SG methods in low-data regimes, and surpasses SS methods with few tens of training samples. On average across tasks, PSFNet achieves 3.1 dB higher PSNR, 2.8% higher SSIM, and 0.3 × lower RMSE than baselines. Furthermore, in both supervised and unsupervised setups, PSFNet requires an order of magnitude lower samples compared to SG methods, and enables an order of magnitude faster inference compared to SS methods. Thus, the proposed model improves deep MRI reconstruction with elevated learning and computational efficiency.

1. Introduction

The unparalleled soft-tissue contrast and non-invasiveness of MRI render it a preferred modality in many diagnostic applications [1,2], and downstream imaging tasks such as classification [3] and segmentation [4,5]. However, the adverse effects of low spin polarization at mainstream field strengths on the signal-to-noise ratio make it slower against alternate modalities such as CT [6]. Since long scan durations inevitably constrain clinical utility, there is an ever-growing interest in accelerated MRI methods to improve scan efficiency. Accelerated MRI involves an ill-posed inverse problem with the aim of mapping under-sampled acquisitions in k-space to high-quality images corresponding

---

* Corresponding author at: Department of Electrical and Electronics Engineering, Bilkent University, Ankara 06800, Turkey.
E-mail address: cukur@ee.bilkent.edu.tr (T. Çukur).
[1] Equal contribution.

to fully-sampled acquisitions. The traditional framework for solving this problem relies on parallel imaging (PI) capabilities of receive coil arrays [7,8], in conjunction with hand-constructed MRI priors [9,10]. A joint objective is iteratively optimized comprising a data-consistency (DC) term based on the physical signal model, and a regularization term that enforces the MRI prior [9]. The physical model constrains reconstructed data to be consistent with acquired data while considering coil sensitivities and undersampling patterns [11]. Meanwhile, the regularization term, often based on a linear transform where data are assumed to be compressible [9]. Traditional methods typically suffer from high computational burden and suboptimal performance as the distribution of actual MRI data diverges from the distribution parametrized by the hand-constructed MRI prior.

The deep learning (DL) framework based on data-driven priors has recently been adopted to address the limitations of traditional methods [12–16]. DL methods can be broadly categorized into two classes depending on whether they employ scan-specific (SS) or scan-general (SG) priors. SS priors promise improved generalization performance by leveraging the physical signal model to learn a reconstruction specific to each test scan, i.e. undersampled k-space data from a given test subject. Similar to autocalibration procedures in PI, a first group of SS methods perform training using a fully-sampled calibration region and then exercise learned dependencies in broader k-space [15–18]. For instance, [19] proposed a calibration consistency framework to derive scan-specific features from calibration data in central k-space, and integrate these features into the reconstruction model. [20] proposed a multiple-weighting approach to improve the performance of SS priors for high acceleration rates and small calibration regions in k-space. [21] proposed a Bayesian framework and nonlinear inversion to enable the capture of SS priors in the absence of a calibration region. Following the deep image prior technique, a second group of methods use unconditional CNNs as a native MRI prior [22–24]. These CNNs map low-dimensional latent variables onto MR images, and latents and network weights are optimized to ensure consistency to acquired data based on the physical signal model. For example, [25] introduced a generative neural network whose weights were optimized during inference. [26] performed reconstruction via an implicit neural network that mapped low-dimensional latent variables onto MR images. Typically, SS priors are learned on each subject at test time to avoid the need for separate training datasets and offer enhanced reliability in handling atypical anatomy. Nevertheless, they suffer from long inference times that can be prohibitive particularly when deep networks have to be optimized [27–29].

As an alternative DL approach, SG priors can offer more efficient inference by learning latent features of MR images from a dedicated dataset at training time and keeping the learned features frozen at test time [12–14,30–38]. Numerous successful architectures have been reported including perceptrons [39], basic convolutional neural networks (CNNs) [40–43], residual or recurrent CNNs [34,44–46], generative adversarial networks (GANs) [47–51], transformers [22,52] and diffusion models [53,54]. Physics-guided unrolled methods have received particular attention that combine the physical signal model as in traditional frameworks and regularization via a deep network serving as an SG prior [13,32,55–57]. Reconstruction is achieved via serially alternated projections through the physical signal model and the SG prior [43,45, 58–60]. For instance, [36] combined a convolutional recurrent neural network with an optimization algorithm that included the physical signal model. [61] instead used a hybrid deep network composed of convolutional and transformer blocks to improve the capture of long-range context. [62] inserted a quality-control module into an unrolled architecture to guide the network regarding the success of the reconstruction process. [63] introduced a noise-parameter in each cascade of an unrolled architecture to help improve reconstruction performance by adapting to varying noise levels. These SG priors are reported to be performant when trained on relatively large datasets, containing on the order of several thousand image samples or more [33,64,65].

Meeting this data demand can require scanning up to several hundreds of subjects. Unfortunately, extensive scanning proves difficult in many practical scenarios, resulting in training sets that are a few orders of magnitude smaller in size (e.g., containing several tens of samples). Under such low-data regimes, the performance of SG priors can be compromised due to suboptimal training [6,66,67].

To address the drawbacks associated with stand-alone use of SS or SG priors, here we propose a novel parallel-stream fusion model (PSFNet) that consolidates SS and SG priors to enable data-efficient training and computation-efficient inference in deep MRI reconstruction[2]. PSFNet leverages an SS stream to perform linear reconstruction based on the physical signal model, and an SG stream to perform nonlinear reconstruction based on a deep network. Unlike conventional unrolled methods based on serial projections that can cause error accumulation across cascades, PSFNet is based on a parallel stream architecture with learnable fusion of SS and SG priors. Fusion parameters are adapted across cascades and training iterations to emphasize task-critical information. By integrating SS and SG priors in this manner, PSFNet aims to benefit from the short inference times of SG priors and generalization performance of SS priors under low-data regimes. Through these achievements, PSFNet alleviates the dependency on large training sets, improving its practicality. Comprehensive experiments on brain MRI datasets are reported to demonstrate PSFNet under both supervised and unsupervised settings [69–73]. PSFNet is compared against an unrolled SG method [32], two SS methods [17, 74], and conventional SPIRiT reconstructions [11]. Compared to the unrolled model, PSFNet lowers training data requirements an order of magnitude. Compared to SS models, PSFNet offers significantly faster inference times. Our main contributions are:

- A novel cascaded network architecture is introduced that adaptively fuses SS and SG priors across cascades and training iterations to improve learning-based MRI reconstruction in low-data regimes.
- The SS prior facilitates learning of the SG prior with limited data, and empowers PSFNet to successfully generalize to out-of-domain samples.
- The SG prior improves performance by capturing nonlinear residuals, and enhances resilience against suboptimal hyperparameter selection in the SS component.
- Parallel-stream fusion of SS and SG priors yields robust performance with limited training data in both supervised and unsupervised settings.

2. Theory

2.1. Image reconstruction in accelerated MRI

MRI reconstruction is an inverse problem that aims to recover an image from a respective undersampled acquisition:

$$MFx = y \tag{1}$$

where F is the Fourier transform, M is the sampling mask defining acquired k-space locations, x is the multi-coil image to be reconstructed and y are acquired multi-coil k-space data. To improve problem conditioning, additional prior information regarding the expected distribution of MR images is incorporated in the form of a regularization term:

$$\hat{x} = \arg\min_x \quad \lambda \|MFx - y\|_2^2 + R(x) \tag{2}$$

where the first term enforces DC between reconstructed and acquired k-space data, $R(x)$ reflects the MRI prior, and λ controls the balance between the DC and regularization terms.

The DC term can be implemented by injecting the acquired values of k-space data into the reconstruction [13]. Thus, mapping through a DC block is given as:

$$f_{DC}(x) = F^{-1}\Lambda Fx + \frac{\lambda}{1+\lambda}F^{-1}y \tag{3}$$

where Λ is a diagonal matrix with diagonal entries set to $\frac{1}{1+\lambda}$ at acquired k-space locations and set to 1 in unacquired locations.

In traditional methods, the regularization term is based on a hand-constructed transform domain where data are assumed to have a sparse representation [9]. For improved conformation to the distribution of MRI data, recent frameworks instead adopt deep network models to capture either SG priors learned from a large MRI database with hundreds of subjects, or SS priors learned from individual test scans. Learning procedures for the two types of priors are discussed below.

SG priors: In MRI, SG priors are typically adopted to suppress aliasing artifacts in the zero-filled reconstruction (i.e., inverse Fourier transform) of undersampled k-space acquisitions [32]. A deep network model that performs de-aliasing can be learned from a large training dataset of undersampled and corresponding fully-sampled k-space acquisitions, and then employed to implement $R(.)$ in Eq. (2) during inference. The regularization term based on an SG prior is given as:

$$R_{SG}(x) = \arg\min_{x} \| C_{SG}(F^{-1}y;\hat{\theta}_{SG}) - x \|_2^2 \tag{4}$$

where C_{SG} is an image-domain deep network with learned parameters $\hat{\theta}_{SG}$. The formulation in Eq. (4) assumes that C_{SG} recovers multi-coil output images provided multi-coil input images. The parameters θ_{SG} for C_{SG} can be learned based on a pixel-wise loss between reconstructed and ground-truth images. Training is conducted offline via an empirical risk minimization approach based on Monte Carlo sampling [13]:

$$\mathcal{L}_{SG}(\theta_{SG}) = \sum_{n=1}^{N} \| C_{SG}(F^{-1}y^n;\theta_{SG}) - \check{x}^n \|_p \tag{5}$$

where N is the number of training scans, n is the training scan index, $\|.\|_p$ denotes ℓ_p norm, $\check{x}^n$ is the ground-truth multi-coil image derived from the fully-sampled acquisition for the nth scan, and y^n are respective undersampled k-space data.

A common approach to build C_{SG} is based on unrolled architectures that perform cascaded projections through CNN blocks to regularize the image and DC blocks to ensure conformance to the physical signal model [32]. Given a total of K cascades with tied CNN parameters across cascades, the mapping through the kth cascade is [13,75,76]:

$$x_k^r = f_{DC}\left(f_{SG}\left(x_{k-1}^r;\theta_{SG}\right)\right) \tag{6}$$

where x_k^r is the image for the rth scan (that could be a training or test scan) at the output of the kth cascade ($k \in [1,2,\ldots,K]$), and $x_0^r = F^{-1}y^r$ where y^r are the acquired undersampled data for the rth scan. Meanwhile, f_{SG} is the CNN block embedded in the kth cascade with parameters θ_{SG}.

As the parameters of SG priors are trained offline and then frozen during inference, deeper network architectures can be used for enhanced reconstruction performance along with fast inference. However, learning deep networks requires substantial training datasets that may be difficult to collect. Moreover, since SG priors learn aggregate representations of MRI data across training subjects, they may show poor generalization to subject-specific variability in anatomy [22].

SS priors: Unlike SG priors, SS priors are not learned from a dedicated training dataset but instead they are learned directly for individual test scans to improve generalization [15]. The SS prior can also be used to implement $R(.)$ in Eq. (2) with the respective regularization term expressed as:

$$R_{SS}(x) = \arg\min_{x} \| C_{SS}(F^{-1}y;\hat{\theta}_{SS}) - x \|_2^2 \tag{7}$$

where C_{SS} is an image-domain network with parameters $\hat{\theta}_{SS}$. In the absence of ground-truth images, the parameters θ_{SS}^q for the qth test scan can be learned based on proxy k-space losses between reconstructed and acquired undersampled data [27]. Learning is conducted online to minimize this proxy loss:

$$\mathcal{L}_{SS}(\theta_{SS}^q) = \| MFC_{SS}(F^{-1}y^q;\theta_{SS}^q) - y^q \|_p \tag{8}$$

where y^q are acquired undersampled k-space data for the qth scan. An unrolled architecture can be adopted to build C_{SS} by performing cascaded projections through network and DC blocks, resulting in the following mapping for the kth cascade:

$$x_k^q = f_{DC}\left(f_{SS}\left(x_{k-1}^q;\theta_{SS}^q\right)\right) \tag{9}$$

f_{SS} can be operationalized as a linear or nonlinear network [27,28]. As the parameters of SS priors are learned independently for each test scan, they promise enhanced generalization to subject-specific anatomy. However, since training is performed online during inference, SS priors can introduce substantial computational burden, particularly when deep nonlinear networks are used that also increase the risk of overfitting [78].

2.2. PSFNet

Here, we propose to combine SS and SG priors to maintain a favorable trade-off between generalization performance and computational efficiency under low-data regimes. In the conventional unrolling framework, this requires computation of serially alternated projections through the SS, SG and DC blocks:

$$x_k^r = f_{DC}\left(f_{SG}\left(f_{SS}\left(x_{k-1}^r;\theta_{SS}^r\right);\theta_{SG}\right)\right) \tag{10}$$

The unrolled architecture with K cascades can be learned offline using the training set. Note that scarcely-trained SG blocks under low-data regimes can perform suboptimally, introducing residual errors in their output. In turn, these errors will accumulate across serial projections to degrade the overall performance.

To address this limitation, here we introduce a novel architecture, PSFNet, that performs parallel-stream fusion of SS and SG priors as opposed to the serial combination in conventional unrolled methods. PSFNet utilizes a nonlinear SG prior for high performance, and a linear SS prior to enhance generalization without excessive computational burden. The two priors undergo parallel-stream fusion with learnable fusion parameters η and γ, as displayed in Fig. 1. These parameters adaptively control the relative weighting of information extracted by the SG versus SS streams during the course of training in order to alleviate error accumulation. As such, the mapping through the kth cascade in PSFNet is:

$$x_k^r = \eta_k f_{DC}(f_{SS}(x_{k-1}^r;\theta_{SS}^r)) + \gamma_k f_{DC}(f_{SG}(x_{k-1}^r;\theta_{SG})) \tag{11}$$

In Eq. (11), the learnable fusion parameters for the SS and SG blocks at the kth cascade are η_k and γ_k, respectively. To enforce fidelity to acquired data, DC projections are performed on the outputs of SG and SS blocks.

The parallel-stream topology in PSFNet serves as the architectural substrate for segregating SG and SS priors, where each prior has its own branch as shown in Fig. 1. A fundamental distinction between these priors is the way that their parameters are learned. In PSFNet, learning is performed via a nested optimization procedure with subproblems for SG and SS priors. In the outer subproblem, parameters of the SG prior are learned via optimization across all samples in the training set, and these parameters are frozen during inference on the test set. In the inner subproblem, parameters of the SS prior are learned via independent optimization for each individual sample in the training and test sets, so these parameters are not frozen during inference. This nested optimization enables segregation of SG and SS priors, as detailed in the training and inference procedures described below.

a) Parallel-Stream Fusion Model (PSFNet)
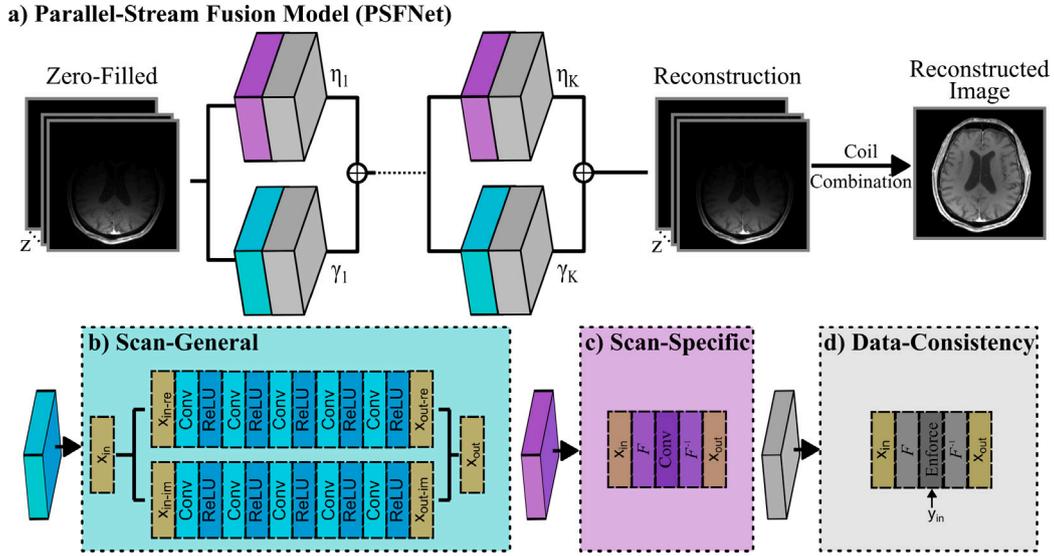
Fig. 1. (a) PSFNet comprises a parallel-stream cascade of sub-networks where each sub-network contains (b) a scan-general (SG) block, and (c) a scan-specific (SS) block. The two parallel blocks are each succeeded by (d) a data-consistency (DC) block, and their outputs are aggregated with learnable fusion weights, η_k and γ_k where k is the cascade index. At the end of K cascades, coil-combination is performed on multi-coil data using sensitivity maps estimated via ESPIRiT [77]. The SG block is implemented as a deep convolutional neural network (CNN) and the SS block was implemented as a linear projection layer.

Training: PSFNet involves a training phase to learn model parameters for the SG prior as well as its fusion with the SS prior. For each individual scan in the training set, PSFNet learns a dedicated SS prior for the given scan. Since learning of a nonlinear SS prior has substantial computational burden, we adopt a linear SS prior in PSFNet. In particular, the SS block performs dealiasing via convolution with a linear kernel [77]:

$$f_{SS}(x_{k-1}^n; \theta_{SS}^n) = F^{-1}\{\theta_{SS}^n \circledast F x_{k-1}^n\} \tag{12}$$

where $\theta_{SS}^n \in \mathbb{C}^{(z \times z \times w \times w)}$ with n denoting the training scan index, z denoting the number of coil elements, and w denoting the kernel size in k-space. The SS blocks contain unlearned Fourier and inverse Fourier transformation layers as their input and output layers, respectively, and convolution is computed over the spatial frequency dimensions in k-space. Meanwhile, the SG prior is implemented as a deep CNN operating in image domain:

$$f_{SG}(x_{k-1}^n; \theta_{SG}) = CNN(x_{k-1}^n) \tag{13}$$

Across the scans in the training set, the training loss for PSFNet can then be expressed in constrained form as:

$$\mathcal{L}_{PSFNet}(\theta_{SG}, \boldsymbol{\gamma}, \boldsymbol{\eta}) = \sum_{n=1}^{N} \| \eta_K f_{DC}(f_{SS}(x_{K-1}^n; \hat{\theta}_{SS}^n)) $$
$$+ \gamma_K f_{DC}(f_{SG}(x_{K-1}^n; \theta_{SG})) - \check{x}^n \|_p$$
$$\text{s.t. } \hat{\theta}_{SS}^n = \underset{\theta_{SS}^n}{\arg\min} \| F^{-1} W^n y^n - f_{SS}(F^{-1} W^n y^n; \theta_{SS}^n) \|_2^2 \tag{14}$$

The constraint in Eq. (14) corresponds to the inner optimization sub-problem to learn the parameters of the SS prior per individual samples $\hat{\theta}_{SS}^n$. The learned parameters of the SS prior are then used to solve the outer subproblem for the SG prior across samples in the training set. Assuming that the linear relationships among neighboring spatial frequencies are similarly distributed across k-space [77], $\hat{\theta}_{SS}^n$ is learned by solving a self-regression problem on the subset of fully-sampled data in central k-space, where W^n is a mask operator that selects data within this calibration region.

Note that, unlike deep reconstruction models purely based on SG priors, the SG prior in PSFNet is not directly trained to remove artifacts in zero-filled reconstructions of undersampled data. Instead, the SG prior is trained to concurrently suppress artifacts in reconstructed images along with the SS prior; and the relative importance attributed to

the two priors is determined by the fusion parameters at each cascade. As such, the SS prior can be given higher weight during initial training iterations where the SG prior is scarcely trained, whereas its weight can be relatively reduced during later iterations once the SG prior has been sufficiently trained. This adaptive fusion approach thereby lowers reliance on the availability of large training sets.

Inference: During inference on the qth test scan, the respective SS prior is learned online as:

$$\hat{\theta}_{SS}^q = \underset{\theta_{SS}^q}{\arg\min} \| F^{-1} W^q y^q - f_{SS}^q(F^{-1} W^q y^q; \theta_{SS}^q) \|_2^2 \tag{15}$$

Afterwards, the learned $\hat{\theta}_{SS}^q$ is used along with the previously trained $\hat{\theta}_{SG}$ to perform repeated projections through K cascades as described in Eq. (11). The multi-coil image recovered by PSFNet at the output of the K cascade is:

$$\hat{x}^q = \eta_K f_{DC}(f_{SS}(x_{K-1}^q; \hat{\theta}_{SS}^q)) + \gamma_K f_{DC}(f_{SG}(x_{K-1}^q; \hat{\theta}_{SG})) \tag{16}$$

where $\hat{x}^q$ denotes the recovered image. The final reconstruction can be obtained by performing combination across coils:

$$\hat{x}_{combined}^q = A^* \hat{x}^q \tag{17}$$

where A are coil sensitivities, and A^* denotes the conjugate of A.

3. Methods

3.1. Implementation details

In each cascade, PSFNet contained two parallel streams with SG and SS blocks. The SG blocks comprised an input layer followed by a stack of 4 convolutional layers with 64 channels and 3×3 kernel size each, and an output layer with ReLU activation functions. They processed complex images with separate channels for real and imaginary components. The SS blocks comprised a Fourier layer, 5 projection layers with identity activation functions, and an inverse Fourier layer. They processed complex images directly without splitting real and imaginary components. The linear convolution kernel used in the projection layers was learned from the calibration region by solving a Tikhonov regularized self-regression problem [11]. The DC blocks comprised 3 layers respectively to implement forward Fourier transformation, restoration of acquired k-space data and inverse Fourier

transformation. PSFNet was implemented with 5 cascades, $K = 5$. The weights of SG, SS, and DC blocks were tied across cascades to limit model complexity [32]. The only exception were fusion coefficients that determine the relative weighting of the SG and SS blocks at each stage ($\gamma_1, \ldots, \gamma_k, \ldots, \gamma_5$ $\eta_1, \ldots \eta_k, \ldots, \eta_5$). These fusion parameters were kept distinct across cascades. Coil-combination on the recovered multi-coil images was performed using sensitivity maps estimated via ESPIRiT [77].

3.2. MRI dataset

Experimental demonstrations were performed using brain MRI scans from the NYU fastMRI database [79]. Here, contrast-enhanced T_1-weighted (cT_1-weighted) and T_2-weighted acquisitions were considered. The fastMRI dataset contains volumetric MRI data with varying image and coil dimensionality across subjects. Note that a central aim of this work was to systematically examine the learning capabilities of models for varying number of training samples. To minimize potential biases due to across-subject variability in MRI protocols, here we selected subjects with matching imaging matrix size and number of coils. To do this, we only selected subjects with at least 10 cross-sections and only the central 10 cross-sections were retained in each subject. We further selected subjects with an in-plane matrix size of 256×320 for cT_1 acquisitions, and of 288×384 for T_2 acquisitions. Background regions in MRI data with higher dimensions were cropped. Lastly, we restricted our sample selection to subjects with at least 5 coil elements, and geometric coil compression [80] was applied. Experiments were conducted following compression onto 5 or 3 virtual coils.

Fully-sampled acquisitions were retrospectively undersampled to achieve acceleration rates of R = 4x, 6x and 8x. Random undersampling patterns were designed via either a bi-variate normal density function peaking at the center of k-space, or a uniform density function across k-space. The standard deviation of the normal density function was adjusted to maintain the expected value of R across k-space. The fully-sampled calibration region spanned a 40×40 window in central k-space.

3.3. Competing methods

PSFNet was compared against several state-of-the-art approaches including SG methods, SS methods, and traditional PI reconstructions. For methods containing SG priors, both supervised and unsupervised variants were implemented.

PSFNet: A supervised variant of PSFNet was trained using paired sets of undersampled and fully-sampled acquisitions.

PSFNet$_{US}$: An unsupervised variant of PSFNet was implemented using self-supervision based on only undersampled training data. Acquired data were split into two non-overlapping sets where 40% of samples was reserved for evaluating the training loss and 60% of samples was reserved to enforce DC [71].

MoDL: A supervised SG methods based on an unrolled architecture with tied weights across cascades was used [32]. MoDL serially interleaves SG and DC blocks. The number of cascades and the structure of SG and DC blocks were identical to those in PSFNet.

MoDL$_{US}$: An unsupervised variant of MoDL was implemented using self-supervision. A 40%–60% split was performed on acquired data to evaluate the training loss and enforce data consistency, respectively [71].

sRAKI-RNN: An SS method was implemented based on the MoDL architecture [74]. Learning was performed to minimize DC loss on the fully-sampled calibration region. Calibration data were randomly split with 75% of samples used to define the training loss and 25% of samples reserved to enforce DC. Multiple input–output pairs were produced for a single test sample by utilizing this split.

SPIRiT: A traditional PI reconstruction was performed using the SPIRiT method [11]. Reconstruction parameters including the regularization weight for kernel estimation (κ), kernel size (w), and the number of iterations (N_{iter}) were independently optimized for each reconstruction task via cross-validation.

SPARK: An SS method was used to correct residual errors from an initial SPIRiT reconstruction [17]. Learning was performed to minimize DC loss on the calibration region. The learned SS prior was then used to correct residual errors in the remainder of the k-space.

rGAN: A conditional GAN model that used adversarial learning was implemented using the architecture, loss functions, and hyperparameters outlined in [48].

DDPM: An unconditional diffusion model was implemented using the architecture, loss functions, and hyperparameters outlined in [81]. During inference, reverse diffusion steps were interleaved with data-consistency projections.

3.4. Optimization procedures

For all methods, hyperparameter selection was performed via cross-validation on a three-way split of data across subjects. There was no overlap among training, validation and test sets in terms of subjects. Data from 10 subjects were reserved for validation, and data from a separate set of 40 subjects were reserved for testing. The number of subjects in the training set was varied from 1 to 50. Hyperparameters that maximized peak signal-to-noise ratio (PSNR) on the validation set were selected for each method.

Training was performed via the Adam optimizer with learning rate $\zeta = 10^{-4}$, $\beta_1 = 0.90$ and $\beta_2 = 0.99$ [82]. All deep learning methods were trained to minimize hybrid $\ell_1 - \ell_2$-norm loss between recovered and target data (e.g., between reconstructed and ground truth images for PSFNet, between recovered and acquired k-space samples for PSFNet$_{US}$) [71]. For PSFNet and MoDL, the selected number of epochs was 200, batch size was set to 2 for the limited number of training samples ($N_{samples} < 10$), and to 5 otherwise. In DC blocks, $\lambda = \infty$ was used to enforce strict data consistency. For PSFNet and SPIRiT, the kernel width (w) and regularization parameter (κ) values were set as (κ, w) = (10^{-2}, 9) at R = 4 and (10^{-2}, 9) at R = 8 for cT_1-weighted reconstructions, and as (10^0, 17) at R = 4 and (10^{-2}, 17) at R = 8 for T_2-weighted reconstructions. For SPIRiT, the number of iterations N_{iter} was set as 13 at R = 4 and 27 at R = 8 for cT_1-weighted reconstructions, 20 at R = 4 and 38 at R = 8 for T_2-weighted reconstructions. For sRAKI-RNN, the selected number of epochs was 500 and batch size was set to 32. All other optimization procedures were identical to MoDL. For SPARK, network architecture and training procedures were adopted from [17], except for the number of epochs (N_{epoch}) and learning rate (ζ) which were optimized on the validation set as (N_{epoch}, ζ)= (100, 10^{-2}) For cT_1-weighted reconstructions, and (N_{epoch}, ζ) = (250, 10^{-3}) for T_2-weighted reconstructions.

All competing methods were executed on an NVidia RTX 3090 GPU, and models were coded in Tensorflow except for SPARK which was implemented in PyTorch. SPARK was implemented using the toolbox at https://github.com/YaminArefeen/spark_mrm_2021. The code to implement PSFNet will be available publicly at https://github.com/icon-lab/PSFNet upon publication.

3.5. Performance metrics

Performance assessments for reconstruction methods were carried out by visual observations and quantitative metrics. PSNR and structural similarity index (SSIM) were used for quantitative evaluation. For each method, metrics were computed on coil-combined images from the reconstruction and from the fully-sampled ground truth acquisition. Statistical differences between competing methods were examined via non-parametric Wilcoxon signed-rank tests.

3.6. Experiments

Several different experiments were conducted to systematically examine the performance of competing methods. Assessments aimed

to investigate reconstruction performance under low training data regimes, generalization performance in case of mismatch between training and testing domains, contribution of the parallel-stream design to reconstruction performance, sensitivity to hyperparameter selection, performance in unsupervised learning, and computational complexity.

Performance in low-data regimes: Deep SG methods for MRI reconstruction typically suffer from suboptimal performance as the size of the training dataset is constrained. To systematically examine reconstruction performance, we trained supervised variants of PSFNet and MoDL while the number of training samples ($N_{samples}$) was varied in the range [2–500] cross sections. To attain a given number of samples, sequential selection was performed across subjects and across cross-sections within each subject. Thus, the number of unique subjects included in the training set roughly corresponded to $N_{samples}/10$ (since there were 10 cross-sections per subject). SS reconstructions were also performed with sRAKI-RNN, SPIRiT and SPARK. In the absence of fully-sampled ground truth data to guide the learning of the prior, unsupervised training of deep reconstruction models may prove relatively more difficult compared to supervised training. In turn, this may elevate requirements on training datasets for unsupervised models. To examine data efficiency for unsupervised training, we compared the reconstruction performance of PSFNet$_{US}$ and MoDL$_{US}$ as $N_{samples}$ was varied in the range of [2–500] cross sections. Comparisons were also provided against sRAKI-RNN, SPIRiT and SPARK.

Generalization performance: Deep reconstruction models can suffer from suboptimal generalization when the MRI data distribution shows substantial variation between the training and testing domains. To examine generalizability, PSFNet models were trained on data from a source domain and tested on data from a different target domain. The domain-transferred models were then compared to models trained and tested directly in the target domain. Three different factors were altered to induce domain variation: tissue contrast, undersampling pattern, and acceleration rate. First, the capability to generalize to different tissue contrasts was evaluated. Models were trained on data from a source contrast and tested on data from a different target contrast. Domain-transferred models were compared to target-domain models trained on data from the target contrast. Next, the capability to generalize to different undersampling patterns was assessed. Models were trained on data undersampled with variable-density patterns and tested on data undersampled with uniform-density patterns. Domain-transferred models were compared to target-domain models trained on uniformly undersampled data. Lastly, the capability to generalize to different acceleration rates was examined. Models were trained on acquisitions accelerated at R = 4x and tested on acquisitions accelerated at R = 8x. Domain-transferred models were compared to target-domain models trained at R = 8x.

Computational complexity: Finally, we assessed the computational complexity of competing methods. For each method, training and inference times were measured for a single subject with 10 cross-sections. Each cross-section had an imaging matrix size of 256×320 and contained data from 5 coils. For all methods including SS priors, hyperparameters optimized for cT$_1$-weighted reconstructions at R = 4 were used.

Ablation studies: To assess the contribution of the parallel-stream design in PSFNet, a conventional unrolled variant of PSFNet was formed, named as PSFNet$_{Serial}$. PSFNet$_{Serial}$ combined the SG and SS priors via serial projections as described in Eq. (10). To assess the contribution of the SS and SG prior in PSFNet, an SS-only variant was formed that ablated the SG prior and an SG-only variant was formed that ablated the SS prior. Modeling procedures were kept identical for fair comparison. We also examined the relative influence of the SS and SG priors in PSFNet as a function of the training set size by observing the fusion weights for the two priors. $N_{samples}$ was varied in the range of [2–500] cross sections.

SS priors might show elevated sensitivity to hyperparameter selection as they are learned from individual test scans. We assessed

the reliability of reconstruction performance against suboptimal hyperparameter selection for SS priors. For this purpose, analyses were conducted on SPIRiT, SPARK and PSFNet that embody SS methods to perform linear reconstructions in k-space. The set of hyperparameters examined included regularization parameters for kernel estimation (κ) and kernel size (w). Separate models were trained using κ in range [10^{-3}-10^0] and w in range [5–17].

4. Results

4.1. Performance in low-data regimes

Common SG methods for MRI reconstruction are based on deep networks that require copious amounts of training data, so performance can substantially decline on limited training sets [33,64]. In contrast, PSFNet leverages an SG prior to concurrently reconstruct an image along with an SS prior. Therefore, we reasoned that its performance should scale favorably under low-data regimes compared to SG methods. We also reasoned that PSFNet should yield elevated performance compared to SS methods due to residual corrections from its SG prior. To test these predictions, we trained supervised variants of PSFNet and MoDL along with SPIRiT, sRAKI-RNN, and SPARK while the number of training samples ($N_{samples}$) was systematically varied. Fig. 2 displays average PSNR, SSIM and RMSE performance across cT$_1$- and T$_2$-weighted reconstructions as a function of $N_{samples}$. PSFNet outperforms the scan-general MoDL method for all values of $N_{samples}$ ($p < 0.05$). As expected, performance benefits with PSFNet become more prominent towards lower values of $N_{samples}$. PSFNet also outperforms traditional SPIRiT and scan-specific sRAKI-RNN and SPARK methods broadly across the examined range of $N_{samples}$ ($p < 0.05$). Note that while MoDL requires $N_{samples} = 30$ (3 subjects) to offer on par performance to SS methods, PSFNet yields superior performance with as few as $N_{samples} = 2$. Representative reconstructions for cT$_1$- and T$_2$-weighted images are depicted in Figs. 3 and 4, where $N_{samples} = 10$ from a single subject were used for training. PSFNet yields lower reconstruction errors compared to all other methods in this low-data regime, where competing methods either show elevated noise or blurring.

We then questioned whether the performance benefits of PSFNet are also apparent during unsupervised training of deep network models. For this purpose, unsupervised variants PSFNet$_{US}$ and MoDL$_{US}$ were trained via self-supervision [71]. PSFNet$_{US}$ was compared against MoDL$_{US}$, SPIRiT, sRAKI-RNN, and SPARK while the number of training samples ($N_{samples}$) was systematically varied. Fig. 5 displays average PSNR, SSIM and RMSE performance across cT$_1$- and T$_2$-weighted reconstructions as a function of $N_{samples}$. Similar to the supervised setting, PSFNet$_{US}$ outperforms MoDL$_{US}$ for all values of $N_{samples}$ ($p < 0.05$), and the performance benefits are more noticeable at lower $N_{samples}$. In this case, however, MoDL$_{US}$ is unable to reach the performance of the best performing SS method (SPARK) even at $N_{samples} = 500$. In contrast, PSFNet$_{US}$ starts outperforming SPARK with approximately $N_{samples} = 40$ (4 subjects). The enhanced reconstruction quality with PSFNet$_{US}$ is corroborated in representative reconstructions for cT$_1$- and T$_2$-weighted images depicted in Figs. 6 and 7, where $N_{samples} = 100$ were used for training. Taken together, these results indicate that the data-efficient nature of PSFNet facilitates the training of both supervised and unsupervised MRI reconstruction models.

4.2. Generalization performance

An important advantage of SS priors is that they allow model adaptation to individual test samples, thereby promise enhanced performance in out-of-domain reconstructions [27]. Yet, SG priors with fixed parameters might show relatively limited generalizability during inference [28,83]. To assess generalization performance, we introduced domain variations by altering three experimental factors:

Fig. 2. Average (a) PSNR, (b) SSIM and (c) RMSE across test subjects for cT_1- and T_2-weighted image reconstructions at R = 4x. Model training was performed for varying number of training samples ($N_{samples}$, lower x-axis) and thereby training subjects ($N_{subjects}$, upper x-axis). Results are shown for SPIRiT, SPARK, sRAKI-RNN, MoDL and PSFNet.

Table 1

Generalization across tissue contrasts. PSNR, SSIM, RMSE values (mean±s.e.) across test subjects. Results are shown for scan-specific models (SPIRiT, SPARK, sRAKI-RNN), target-domain models (MoDL, PSFNet) and domain-transferred models (MoDL-DT, PSFNet-DT) at R = 4x, 5 coils. The tissue contrast in the target domain is listed in the left-most column (cT_1 or T_2), domain-transferred models were trained for the non-target tissue contrast.

	SPIRiT	SPARK	sRAKI-RNN	MoDL	MoDL-DT	PSFNet	PSFNet-DT
	PSNR						
cT_1	37.6 ±1.5	37.6 ±1.5	36.8 ±1.3	38.5 ± 1.5	38.2 ±1.5	39.9 ±1.7	39.4 ±1.6
T_2	35.8 ±1.0	36.5 ±1.0	35.2 ±1.1	37.9 ± 1.0	37.5 ±1.1	39.0 ±1.0	39.0 ±0.9
	SSIM						
cT_1	93.1 ±1.5	93.3 ±1.4	93.8 ±1.0	95.1 ±1.0	94.8 ±1.1	95.8 ±1.0	95.6 ±1.0
T_2	90.8 ±1.2	93.1 ±1.0	94.9 ±0.6	96.2 ±0.5	96.2 ±0.5	96.7 ±0.4	96.8 ±0.4
	RMSE (×100)						
cT_1	1.4 ±0.2	1.4 ±0.2	1.5 ±0.2	1.2 ±0.2	1.3 ±0.3	1.0 ±0.2	1.1 ±0.2
T_2	1.6 ±0.2	1.5 ±0.2	1.9 ±0.2	1.3 ±0.2	1.5 ±0.4	1.1 ±0.1	1.2 ±0.2

Table 2

Generalization across tissue contrasts. PSNR, SSIM, RMSE values for competing methods at R = 6x, 5 coils.

	SPIRiT	SPARK	sRAKI-RNN	MoDL	MoDL-DT	PSFNet	PSFNet-DT
	PSNR						
cT_1	35.8 ±1.4	35.8 ±1.4	35.4 ±1.5	36.7 ± 1.5	36.1 ±1.4	37.8 ±1.2	37.4 ±1.2
T_2	34.8 ±1.0	34.8 ±1.0	33.6 ±0.9	35.7 ± 1.1	35.4 ±1.1	37.0 ±1.1	36.6 ±1.1
	SSIM						
cT_1	91.6 ±1.6	91.7 ±1.6	92.5 ±1.2	93.4 ±1.3	93.1 ±1.3	94.4 ±1.2	94.2 ±1.2
T_2	90.1 ±1.3	91.0 ±1.1	93.7 ±0.7	94.7 ±0.6	94.6 ±0.6	95.5 ±0.6	95.3 ±0.6
	RMSE (×100)						
cT_1	1.1 ±0.2	1.1 ±0.2	1.1 ±0.2	1.0 ±0.2	1.0 ±0.2	0.9 ±0.1	0.9 ±0.1
T_2	1.2 ±0.2	1.2 ±0.2	1.3 ±0.2	1.0 ±0.2	1.1 ±0.2	0.9 ±0.1	1.0 ±0.1

tissue contrast, undersampling pattern, and acceleration rate. For methods comprising SG components, we built both target-domain models that were trained in the target domain, and domain-transferred models that were trained in a non-target domain. We then compared the reconstruction performances of the two models in the target domain.

First, we examined generalization performance when the tissue contrast varied between training and testing domains (e.g., trained on cT_1, tested on T_2). Analyses were conducted for $N_{samples} = 500$, and varying configurations of acceleration rate and number of coils. Performance metrics for competing methods are listed for R = 4x, 5 coils in Table 1, for R = 6x, 5 coils in Table 2, and for R = 4x, 3 coils in Table 3. Performance losses incurred during domain transfer for methods based on SG priors are rather modest. Compared to PSFNet, PSFNet-DT shows a slight loss of 0.3 dB PSNR, 0.1% SSIM and 0.1 × RMSE ($p < 0.05$). Similarly, compared to MoDL, MoDL-DT shows a loss of 0.3 dB PSNR, 0.2% SSIM and 0.1 × RMSE ($p < 0.05$). Note that PSFNet-DT is the top-contender for the best performing PSFNet, and it outperforms the closest competing SS method by 2.2 dB PSNR, 2.8% SSIM and 0.4 × RMSE ($p < 0.05$).

Second, we examined generalization performance when models were trained with variable-density and tested on uniform-density undersampling patterns. Table 4 lists performance metrics for competing methods. Compared to MoDL, MoDL-DT suffers a notable performance loss of 3.6 dB PSNR, 2.5% SSIM and 0.3 × RMSE ($p < 0.05$). Meanwhile, compared to PSFNet, PSFNet-DT shows a relatively limited loss of 0.4 dB PSNR, 0.2% SSIM and 0.5 × RMSE ($p < 0.05$). Note that PSFNet-DT is again the second-best method after PSFNet, and it outperforms the closest competing SS method by 3.4 dB PSNR, 3.7% SSIM and 0.3 × RMSE ($p < 0.05$).

Third, we examined generalization performance when models were trained at R = 4x and tested on R = 8x. Table 5 lists performance metrics for competing methods. Compared to MoDL, MoDL-DT suffers a notable performance loss of 1.0 dB PSNR ($p < 0.05$), on par RMSE and SSIM. Compared to PSFNet, PSFNet-DT shows a lower loss of 0.6 dB PSNR and 0.1 × RMSE ($p < 0.05$) and on par SSIM. Note that PSFNet-DT is the second-best method following PSFNet, and it outperforms the closest competing SS method by 1.2 dB PSNR, 1.9% SSIM and 0.2 × RMSE ($p < 0.05$). Taken together, these results clearly suggest that the SS prior in PSFNet contributes to its improved generalization performance over the scan-general MoDL method, while the SG prior in PSFNet enables it to outperform competing SS methods.
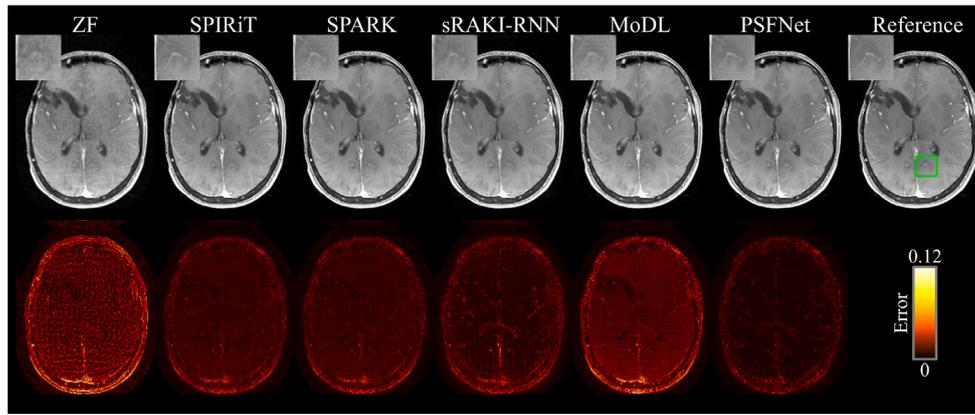
Fig. 3. cT_1-weighted image reconstructions at R = 4x via SPIRiT, SPARK, sRAKI-RNN, MoDL, and PSFNet along with the zero-filled reconstruction (ZF) and the reference image obtained from the fully-sampled acquisition. Error maps for each method are shown in the bottom row. Zoom-in windows display the region annotated with the green box. MoDL and PSFNet were trained on 10 cross-sections from a single subject. SPIRiT, SPARK and sRAKI-RNN directly performed inference on test data without a priori model training. PSFNet shows superior performance to competing methods in terms of residual reconstruction errors.
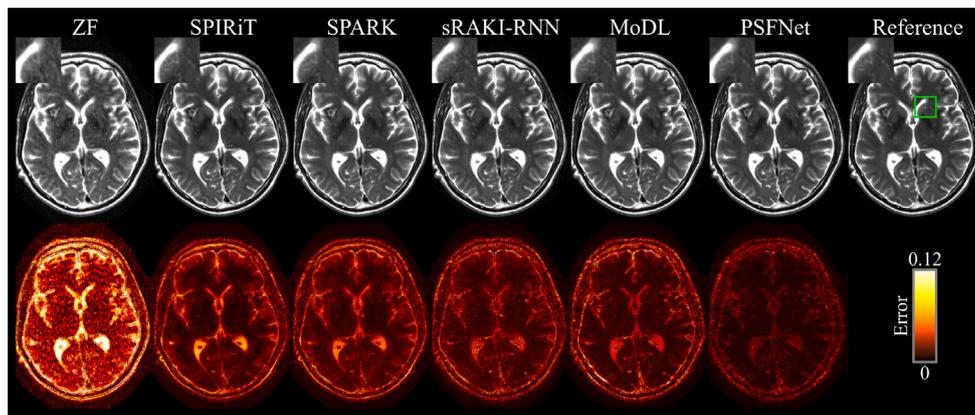
Fig. 4. T_2-weighted image reconstructions at R = 4x via SPIRiT, SPARK, sRAKI-RNN, MoDL, and PSFNet along with the zero-filled reconstruction (ZF) and the reference image obtained from the fully-sampled acquisition. Error maps for each method are shown in the bottom row. MoDL and PSFNet were trained on 10 cross-sections from a single subject. SPIRiT, SPARK and sRAKI-RNN directly performed inference on test data without a priori model training. PSFNet shows superior performance to competing methods in terms of residual reconstruction errors.

Table 3

Generalization across tissue contrasts. PSNR, SSIM, RMSE values for competing methods at R = 4x, 3 coils.

	SPIRiT	SPARK	sRAKI-RNN	MoDL	MoDL-DT	PSFNet	PSFNet-DT
				PSNR			
cT_1	35.1 ±2.8	35.1 ±2.8	35.8 ±1.6	36.8 ± 1.8	36.5 ±1.8	37.9 ±1.9	37.6 ±1.9
T_2	35.7 ±1.3	35.7 ±1.3	34.5 ±1.2	36.4 ± 1.2	36.1 ±1.3	37.6 ±1.3	37.4 ±1.3
				SSIM			
cT_1	91.6 ±4.3	91.6 ±4.3	92.9 ±1.1	94.5 ±1.1	94.1 ±1.2	95.2 ±1.1	95.1 ±1.1
T_2	94.0 ±1.3	94.0 ±1.3	94.2 ±0.7	95.5 ±0.6	95.5 ±0.6	96.2 ±0.5	96.1 ±0.5
				RMSE (×100)			
cT_1	2.0 ±1.3	2.0 ±1.3	1.7 ±0.3	1.5 ±0.3	1.6 ±0.3	1.3 ±0.3	1.4 ±0.3
T_2	1.7 ±0.3	1.7 ±0.3	2.0 ±0.4	1.6 ±0.2	1.6 ±0.3	1.4 ±0.2	1.4 ±0.2

Table 4

Generalization across undersampling patterns. PSNR, SSIM, RMSE values for competing methods at R = 4x, 5 coils. Domain-transferred models were trained with variable-density undersampling, and tested on uniform-density undersampling. Target-domain models were trained and tested with uniform-density undersampling.

	SPIRiT	SPARK	sRAKI-RNN	MoDL	MoDL-DT	PSFNet	PSFNet-DT
				PSNR			
cT_1	37.1 ±1.8	37.1 ±1.7	33.6 ±1.4	37.0 ± 1.7	33.6 ±1.8	40.2 ±1.6	39.9 ±1.6
T_2	35.1 ±1.3	35.6 ±1.3	31.6 ±1.5	37.0 ± 1.1	33.2 ±1.2	40.2 ±1.1	39.7 ±1.2
				SSIM			
cT_1	92.9 ±1.5	93.0 ±1.5	91.2 ±1.5	93.4 ±1.3	91.2 ±2.0	95.9 ±1.2	95.6 ±1.2
T_2	90.6 ±1.5	92.1 ±1.5	91.5 ±1.2	95.6 ±0.7	92.7 ±1.1	97.1 ±0.6	96.9 ±0.6
				RMSE (×100)			
cT_1	1.0 ±0.2	1.0 ±0.2	1.4 ±0.3	1.0 ±0.2	1.5 ±0.4	0.7 ±0.1	0.7 ±0.1
T_2	1.6 ±0.2	1.5 ±0.2	1.9 ±0.2	1.3 ±0.2	1.5 ±0.4	1.1 ±0.1	1.2 ±0.2

Table 5
Generalization across acceleration rates. PSNR, SSIM, RMSE values for competing methods with 5 coils. Domain-transferred models were trained R = 4x and tested at R = 8x. Target-domain models were trained and tested at R = 8x.

	SPIRiT	SPARK	sRAKI-RNN	MoDL	MoDL-DT	PSFNet	PSFNet-DT
				PSNR			
cT_1	34.7 ±1.5	34.8 ±1.5	34.3 ±1.5	35.3 ± 1.4	34.5 ±1.7	36.5 ±1.5	36.2 ±1.5
T_2	33.6 ±1.0	33.7 ±1.0	32.6 ±0.9	34.6 ± 1.0	33.4 ±1.2	35.6 ±1.1	34.6 ±1.2
				SSIM			
cT_1	89.8 ±1.9	90.8 ±1.6	91.4 ±1.4	92.1 ±1.5	92.2 ±1.4	93.3 ±1.4	93.3 ±1.4
T_2	89.0 ±1.3	90.1 ±1.1	92.7 ±0.9	93.5 ±0.8	93.7 ±0.8	94.6 ±0.7	94.5 ±0.7
				RMSE (×100)			
cT_1	1.2 ±0.3	1.2 ±0.3	1.2 ±0.3	1.1 ±0.2	1.2 ±0.3	1.0 ±0.2	1.0 ±0.2
T_2	1.3 ±0.2	1.3 ±0.2	1.4 ±0.2	1.2 ±0.2	1.3 ±0.2	1.0 ±0.2	1.2 ±0.2

Table 6
Computational complexity of competing methods. Training and inference times for data from a single subject, with 10 cross-sections, imaging matrix size 256 × 320 and 5 coils.

	SPIRiT	SPARK	sRAKI-RNN	MoDL	PSFNet
Training (s)	–	–	–	132	337
Inference (s)	0.85	23.35	285.00	0.25	1.13

Table 7
Generalization across tissue contrasts. PSNR, SSIM, RMSE values (mean±s.e.) across test subjects. Results are shown for SPARK, sRAKI-RNN, MoDL, rGAN, DDPM and PSFNet at R = 4x, 5 coils.

	SPARK	sRAKI-RNN	MoDL	rGAN	DDPM	PSFNet
			PSNR			
cT_1	37.6 ±1.5	36.8 ±1.3	38.5 ±1.5	33.5 ±1.2	34.8 ±2.1	39.9 ±1.7
T_2	36.5 ±1.0	35.2 ±1.1	37.9 ±1.0	33.8 ±0.9	34.2 ±2.0	39.0 ±1.0
			SSIM			
cT_1	93.3 ±1.4	93.8 ±1.0	95.1 ±1.0	91.5 ±1.3	92.5 ±1.7	95.8 ±1.0
T_2	93.1 ±1.0	94.9 ±0.6	96.2 ±0.5	93.7 ±0.6	94.7 ±0.7	96.7 ±0.4
			RMSE (×100)			
cT_1	1.4 ±0.2	1.5 ±0.2	1.2 ±0.2	1.4 ±0.3	1.3 ±0.4	1.0 ±0.2
T_2	1.5 ±0.2	1.9 ±0.2	1.3 ±0.2	1.3 ±0.2	1.3 ±0.4	1.1 ±0.1

Fig. 5. Average (a) PSNR, (b) SSIM and (c) RMSE across test subjects for cT_1- and T_2-weighted image reconstructions at R = 4x. Model training was performed for varying number of training samples ($N_{samples}$, lower x-axis) and thereby training subjects ($N_{subjects}$, upper x-axis). Results are shown for SPIRiT, SPARK, sRAKI-RNN, $MoDL_{US}$ and $PSFNet_{US}$.

test samples have a high computational burden. Although PSFNet also embodies an SS prior, its uses a relatively lightweight linear prior as opposed to the nonlinear priors in competing SS methods. Therefore, PSFNet benefits from data-efficient learning while maintaining computationally-efficient inference.

4.4. Radiological assessments

Radiological assessments were conducted on reconstructed images to assess the visual quality of tissue depiction in reference to the ground-truth images derived from fully-sampled acquisitions (Fig. 8). On a 5-point Likert scale, we examined the performance of PSFNet against competing methods, SPIRiT, SPARK, sRAKI-RNN, and MoDL. This evaluation was conducted for cT_1- and T_2-weighted images at R = 4x. PSFNet outperforms all competing methods in opinion scores ($p < 0.05$). On average across tasks, PSFNet achieves 4.4 opinion score, whereas the opinion score is 3.0 for SPIRiT, 2.9 for SPARK, 2.6 for sRAKI-RNN, and 3.5 for MoDL.

4.3. Computational complexity

Next, we assessed the computational complexity of competing methods. Table 6 lists the training times of methods with SG priors, MoDL and PSFNet. Note that SS-based methods do not involve a pre-training step. As it involves learning of an SS prior on each training sample, PSFNet yields elevated training time compared to MoDL. In return, it offers enhanced generalization performance and data-efficient learning. Table 6 also lists the inference times of SPIRiT, SPARK, sRAKI-RNN, MoDL and PSFNet. MoDL and PSFNet that employ SG priors with fixed weights during inference offer fast run times. In contrast, SPARK and sRAKI-RNN that involve SS priors learned on individual

Fig. 6. cT_1-weighted image reconstructions at R = 4x via SPIRiT, SPARK, sRAKI-RNN, $MoDL_{US}$, and $PSFNet_{US}$ along with the zero-filled reconstruction (ZF) and the reference image obtained from the fully-sampled acquisition. Error maps for each method are shown in the bottom row. $MoDL_{US}$ and $PSFNet_{US}$ were trained on 100 cross-sections (from 10 subjects). SPIRiT, SPARK and sRAKI-RNN directly performed inference on test data without a priori model training. $PSFNet_{US}$ shows superior performance to competing methods in terms of residual reconstruction errors.
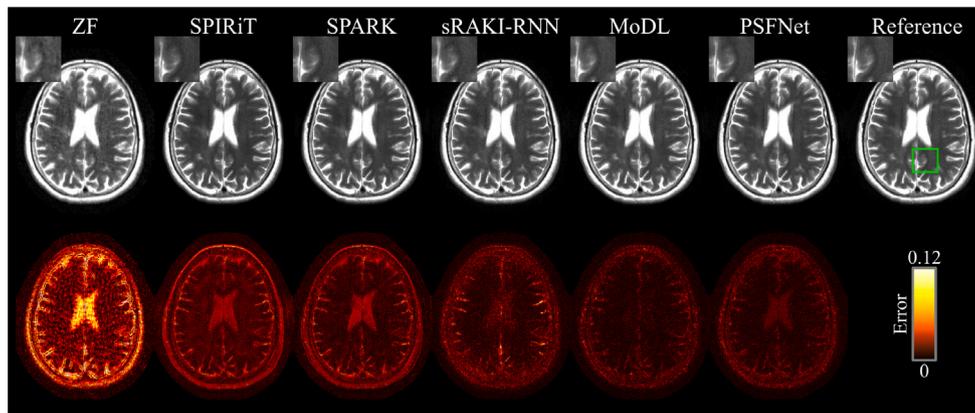
Fig. 7. T_2-weighted image reconstructions at R = 4x via SPIRiT, SPARK, sRAKI-RNN, $MoDL_{US}$, and $PSFNet_{US}$ along with the zero-filled reconstruction (ZF) and the reference image obtained from the fully-sampled acquisition. Error maps for each method are shown in the bottom row. $MoDL_{US}$ and $PSFNet_{US}$ were trained on 100 cross-sections (from 10 subjects). SPIRiT, SPARK and sRAKI-RNN directly performed inference on test data without a priori model training. $PSFNet_{US}$ shows superior performance to competing methods in terms of residual reconstruction errors.
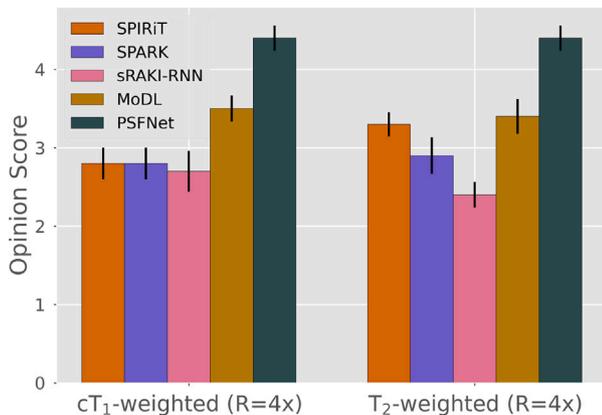
Fig. 8. Radiological opinion scores of SPIRiT, SPARK, sRAKI-RNN, MoDL, and PSFNet on a 5-point Likert scale. Bar plots show mean±s.e. score for cT_1-weighted (left) and T_2-weighted (right) images at R = 4x.

4.5. Comparison to other deep learning methods

We also performed focused comparisons of generalization performance between deep learning methods for MRI reconstruction. For this purpose, PSFNet was compared against SPARK, sRAKI-RNN and MoDL, as well as recent generative methods rGAN [48] and DDPM [81]. Table 7 lists performance metrics for PSFNet and competing methods. PSFNet outperforms all competing methods in each metric ($p < 0.05$). On average over competing methods, PSFNet improves PSNR by 4.5 dB, SSIM by 3.0%, and RMSE by 0.5×10^{-2}. These results suggest that, under low-data regimes, PSFNet not only outperforms SS and SG priors but also recent generative techniques.

4.6. Ablation studies

To demonstrate the value of the parallel-stream fusion strategy in PSFNet over conventional unrolling, PSFNet was compared against a variant model $PSFNet_{Serial}$ that combined SS and SG priors through serially alternated projections. Separate models were trained with number of training samples in the range $N_{samples} = [2\text{–}500]$. Performance in cT_1- and T_2-weighted image reconstruction is displayed in Fig. 9. PSFNet significantly improves reconstruction performance over $PSFNet_{Serial}$ across the entire range of $N_{samples}$ considered ($p < 0.05$), and the benefits grow stronger for smaller training sets. On average across contrasts for $N_{samples} < 10$, PSFNet outperforms $PSFNet_{Serial}$ by 1.8 dB PSNR, 0.6% SSIM and $0.2 \times$ RMSE ($p < 0.05$). These results indicate that the parallel-stream fusion of SG and SS priors in PSFNet is superior to the serial projections in conventional unrolling.

We also examined the contribution of the SS prior and SG prior in PSFNet to reconstruction performance. To do this, an SS-only variant was formed by ablating the SG prior, and an SG-only variant was
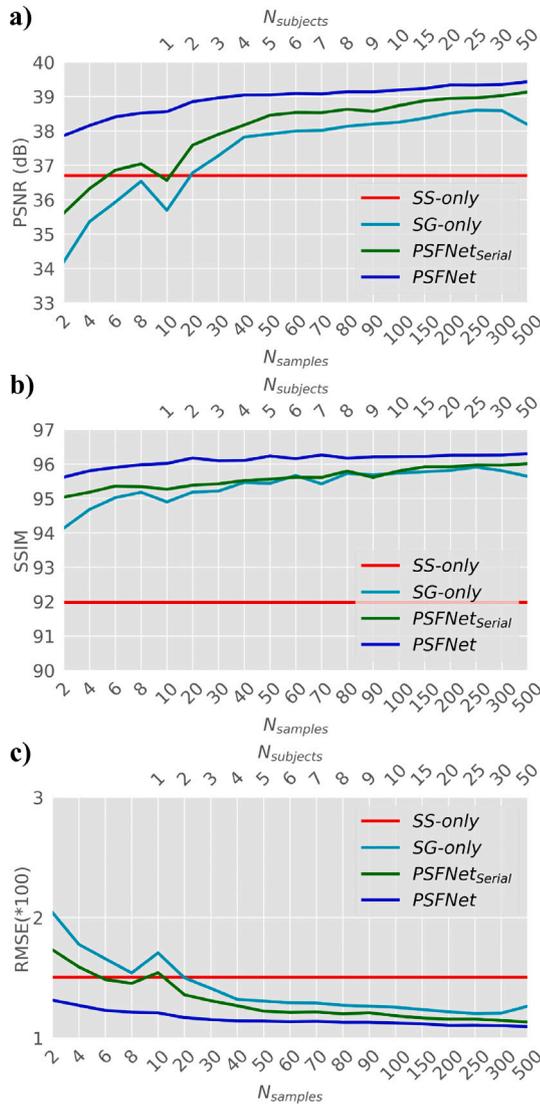
Fig. 9. Average (a) PSNR, (b) SSIM and (c) RMSE values for cT_1- and T_2-weighted image reconstructions at R = 4x. Model training was performed for varying number of training samples ($N_{samples}$, lower x-axis) and thereby training subjects ($N_{subjects}$, upper x-axis). Results are shown for SS-only, SG-only, PSFNet and PSFNet$_{Serial}$.
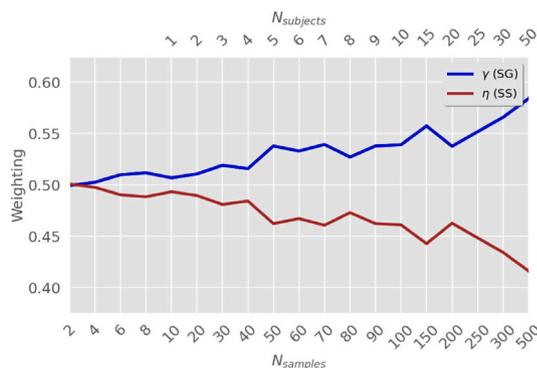
Fig. 10. Weighting of the SG (γ) and SS (η) blocks in the final cascade of PSFNet. Weights were averaged across models trained for cT_1- and T_2-weighted reconstructions at R = 4x. Model training was performed for varying number of training samples ($N_{samples}$, lower x-axis) and thereby training subjects ($N_{subjects}$, upper x-axis). Both blocks are equally weighted with very limited training data. As $N_{samples}$ increases, the weighting of the SG prior becomes more dominant over the weighting of the SS prior.

formed by ablating the SS prior from PSFNet. Separate models were trained with number of training samples in the range $N_{samples} = [2-500]$. Reconstruction performance for PSFNet variants is displayed in Fig. 9. PSFNet significantly improves reconstruction performance over the variants across $N_{samples}$ ($p < 0.05$). On average, PSFNet outperforms the SS-only variant by 2.8 dB PSNR, 4.3% SSIM and $0.5 \times$ RMSE ($p < 0.05$), and the SG-only variant by 1.3 dB PSNR, 0.6% SSIM and $0.2 \times$ RMSE ($p < 0.05$). These results clearly demonstrate that both the SS and SG priors contributed significantly to the performance of PSFNet.

Next, we assessed the relative influence of the SS versus SG priors in the reconstructions computed by PSFNet as a function of training set size. Since the SS prior is independently learned for individual samples, it should not elicit systematic performance variations depending on $N_{samples}$. Thus, we reasoned that the influence of the SG prior should increase towards larger training sets. To test this prediction, we inspected the weightings of the SG (γ) and SS (η) streams as the training set size was varied. Fig. 10 displays weightings at the last cascade as a function of $N_{samples}$. For lower values of $N_{samples}$ where the quality of the SG prior is relatively limited, the SG and SS priors are almost equally weighted. In contrast, as the learning of the SG prior improves with higher $N_{samples}$, the emphasis on the SG prior increases while the SS prior is less heavily weighted.

Parameters of deep networks that implement SS priors are to be learned from a single test sample, so the resultant models can show elevated sensitivity to the selection of hyperparameters compared to SG priors learned from a collection of training samples. Thus, we investigated the sensitivity of PSFNet to key hyperparameters of its SS prior. SPIRiT, SPARK and PSFNet methods all embody a linear k-space reconstruction, so the relevant hyperparameters are the regularization weight and width for the convolution kernel. Performance was evaluated for models were trained in the low-data regime (i.e., $N_{samples} = 10$, 1 subject) for varying hyperparameter values.

Fig. 11 displays PSNR measurements for SPIRiT, SPARK and PSFNet across κ in range (10^{-3}-10^0). While the performance of SPIRiT and SPARK is notably influenced by κ, PSFNet is minimally affected by sub-optimal selection. On average across contrasts, the difference between the maximum and minimum PSNR values is 8.4 dB for SPIRiT, 4.5 dB for SPARK, and a lower 0.7 dB for PSFNet. Note that PSFNet outperforms competing methods across the entire range of κ ($p < 0.05$). Fig. 12 shows PSNR measurements for competing methods across w in range (5–17). In this case, all methods show relatively limited sensitivity to the selection of w. On average across contrasts, the difference between the maximum and minimum PSNR values is 1.5 dB for SPIRiT, 0.5 dB for SPARK, and 0.2 dB for PSFNet. Again, PSFNet outperforms competing methods across the entire range of w ($p < 0.05$). Overall, our results indicate that PSFNet yields improved reliability against sub-optimal hyperparameter selection than competing SS methods.

5. Discussion and conclusion

In this study, we introduced PSFNet for data-efficient training of deep reconstruction models in accelerated MRI. PSFNet is based on synergistic fusion of SS and SG priors in a parallel-stream architecture. The linear SS prior optimizes learning efficiency while maintaining a relatively low computational footprint, whereas the nonlinear SG prior contributes to improved reconstruction performance. In both supervised and unsupervised learning setups, PSFNet notably offers high performance under low-data regimes against baselines. It also maintains competitive inference times to efficient SG methods and reliably generalizes across tissue contrasts, sampling patterns, and acceleration rates. Here PSFNet was primarily for reconstruction of brain images, and it remains important future work to examine its performance in other anatomies.

Several prominent approaches have been introduced in the literature to address the training requirements of deep models based
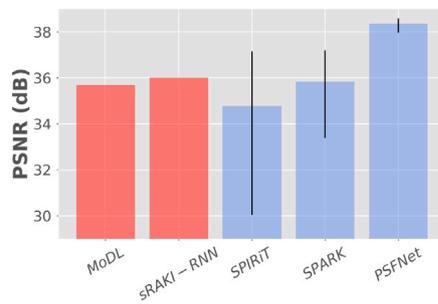
Fig. 11. PSNR measurements were performed on recovered cT$_1$- and T$_2$-weighted images at R = 4x. Bar plots in blue color show average PSNR across $\kappa \in 10^{-3}$-10^1 (i.e., the regularization parameter for kernel estimation). Error bars denote the 90% interval across κ. Bar plots in red color show PSNR for methods that do not depend on the value of κ.
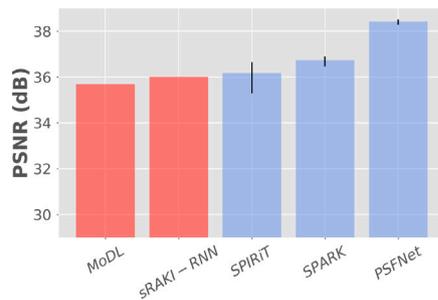
Fig. 12. PSNR measurements were performed on recovered cT$_1$- and T$_2$-weighted images at R = 4x. Bar plots in blue color show the average PSNR across $w \in 5$–17 (i.e., the kernel size). Error bars denote the 90% interval across w. Bar plots in red color show PSNR for methods that do not depend on the value of w.

on SG priors. One approach is to pre-train models on readily available datasets from a separate source domain and then to fine-tune on several tens of samples from the target domain [33,64] or else perform SS fine-tuning [84]. This transfer learning approach relaxes the domain requirements for training datasets. However, the domain-transferred models might be suboptimal when training and testing data distributions are divergent. In such cases, additional training for domain-alignment might be necessary to mitigate performance losses. In contrast, PSFNet contains an SS prior enabling it to better generalize to out-of-domain data without further training. Another approach is to build unsupervised models to alleviate dependency on training datasets with paired undersampled, fully-sampled acquisitions. Model training can be performed either directly on undersampled acquisitions via self-supervision [71] or on unpaired sets of undersampled and fully-sampled acquisitions via cycle-consistent learning [85]. This approach can prove beneficial when fully-sampled acquisitions are costly to collect. Nonetheless, the resulting models still require relatively large datasets form tens of subjects during training [71]. Our experiments on self-supervised variants of PSFNet and MoDL reveal that unsupervised models may require more data compared to their supervised counterparts. As a result, the data-efficiency advantages offered by PSFNet can be especially valuable in the context of unsupervised deep MRI reconstruction.

A fundamentally different framework to lower requirements on training datasets while offering improved generalizability is based on SS priors. In this case, learning can be performed directly on test data and models can be adapted to each scan [15,17]. A group of studies have proposed SS methods based on relatively compact nonlinear models to facilitate learning during inference [15,17,18,86]. However, because learning is performed in central k-space, these methods implicitly assume that local relationships among spatial frequency samples are largely invariant across k-space. While the SS prior in PSFNet also

rests on a similar assumption, the SG components helps correct residual errors that can be introduced due to this assumption. Another group of studies have alternatively adopted the deep image prior (DIP) approach to build SS methods [22,23,27,28]. In DIP, unconditional deep network models that map latent variables onto images are used as native priors for MR images. The priors are learned by ensuring the consistency of reconstructed and acquired data across the entire k-space. Despite improved generalization, these relatively more complex models require increased inference times. In comparison, PSFNet provides faster inference since the weights for its SG prior are fixed, and its SS prior involves a compact linear operator that is easier to learn.

Few independent studies on MRI have proposed approaches related to PSFNet by combining nonlinear and linear reconstructions [6,17,86]. Residual RAKI and SPARK methods initially perform a linear reconstruction, and then use an SS method to correct residual errors via minimizing a DC loss in the calibration region [17,86]. As local relationships among data samples might vary across k-space, the learned SS priors might be suboptimal. Moreover, these methods perform online learning of nonlinear SS priors that introduces relatively high computational burden. In contrast, PSFNet incorporates an SG prior to help improve reliability against sub-optimalities in the SS prior, and uses a linear SS prior for efficiency. Another related method is GrappaNet that improves reconstruction performance by cascading GRAPPA and network-based nonlinear reconstruction steps [6]. While [6] intends to improve image quality, the main aim of our study is to improve practicality by lowering training data requirements of deep models, and improving domain generalizability without elevating inference times. Note that GrappaNet follows the conventional unrolling approach by performing serially alternated projections through linear and nonlinear reconstructions, which can lead to error propagation under low-data regimes [87]. In contrast, PSFNet maintains linear and nonlinear reconstructions as two parallel streams in its architecture, and learns to optimally fuse the information from the two streams.

There are several technical limitations in the current study that can be addressed to further improve performance and reliability. First, PSFNet trains its SG prior with pixel-wise losses than can be amenable to spatial blur. To improve capture of high-frequency information, an adversarial loss term along with a discriminator subnetwork can be included in PSFNet [88]. In this case, an important assessment would be to explore the data-efficiency benefits of PSFNet under adversarial learning. Second, PSFNet implements its SS prior based on linear activation functions for efficiency, albeit this may limit expressiveness. To improve representational capacity, shallow networks with nonlinear activation functions can be considered [86]. While this may increase inference time, it could foster further improvements to generalization performance. Third, the expressiveness of both SS and SG priors might be enhanced by incorporating attention mechanisms as reported in recent studies based on transformer or diffusion models [89,90]. Fourth, here PSFNet is implemented to reconstruct each MRI contrast independently. Improved performance might be viable in multi-contrast MRI protocols by adapting PSFNet to jointly reconstruct and share information across separate MRI contrasts of the same anatomy [60, 91,92]. Fifth, PSFNet primarily lowers demand for training data by incorporating an SS prior. Alternative approaches to reduce reliance on large datasets such as transfer learning can be adopted where SG priors pre-trained on natural images could be employed. It remains important future work to examine the compatibility of the SS prior with these alternative approaches in improving data efficiency. Lastly, there are numerous aspects of MR images including resolution, contrast, noise and artifact levels that can influence their diagnostic utility. For a preliminary assessment, here we conducted a general radiological evaluation on visual quality that implicitly reflects such underlying aspects. More detailed evaluations are warranted in future work to examine the success of reconstruction methods in optimizing individual aspects of MR images.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

References

[1] S. Bauer, R. Wiest, L.-P. Nolte, M. Reyes, A survey of MRI-based medical image analysis for brain tumor studies, Phys. Med. Biol. 58 (13) (2013) R97.

[2] A. Shoeibi, M. Khodatars, M. Jafari, N. Ghassemi, P. Moridian, R. Alizadehsani, S.H. Ling, A. Khosravi, H. Alinejad-Rokny, H. Lam, M. Fuller-Tyszkiewicz, U.R. Acharya, D. Anderson, Y. Zhang, J.M. Gorriz, Diagnosis of brain diseases in fusion of neuroimaging modalities using deep learning: A review, Inf. Fusion 93 (2023) 85–117.

[3] M. Hu, X. Qian, S. Liu, A.J. Koh, K. Sim, X. Jiang, C. Guan, J.H. Zhou, Structural and diffusion MRI based schizophrenia classification using 2D pretrained and 3D naive convolutional neural networks, Schizophr. Res. 243 (2022) 330–341.

[4] K.R.M. Fernando, C.P. Tsokos, Deep and statistical learning in biomedical imaging: State of the art in 3D MRI brain tumor segmentation, Inf. Fusion 92 (2023) 450–465.

[5] Z. Zhu, X. He, G. Qi, Y. Li, B. Cong, Y. Liu, Brain tumor segmentation based on the fusion of deep semantics and edge information in multimodal MRI, Inf. Fusion 91 (2023) 376–387.

[6] A. Sriram, J. Zbontar, T. Murrell, C.L. Zitnick, A. Defazio, D.K. Sodickson, GrappaNet: Combining parallel imaging with deep learning for multi-coil MRI reconstruction, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR, 2020, pp. 14303–14310.

[7] K.P. Pruessmann, M. Weiger, M.B. Scheidegger, P. Boesiger, SENSE: sensitivity encoding for fast MRI, Magn. Reson. Med. 42 (5) (1999) 952–962.

[8] M.A. Griswold, P.M. Jakob, R.M. Heidemann, M. Nittka, V. Jellus, J. Wang, B. Kiefer, A. Haase, Generalized autocalibrating partially parallel acquisitions (GRAPPA), Magn. Reson. Med. 47 (6) (2002) 1202–1210.

[9] M. Lustig, D. Donoho, J.M. Pauly, Sparse MRI: The application of compressed sensing for rapid MR imaging, Magn. Reson. Med. 58 (2007) 1182–1195.

[10] A. Majumdar, Improving synthesis and analysis prior blind compressed sensing with low-rank constraints for dynamic MRI reconstruction, Magn. Reson. Imaging 33 (1) (2015) 174–179.

[11] M. Lustig, J.M. Pauly, SPIRiT: Iterative self-consistent parallel imaging reconstruction from arbitrary k-space, Magn. Reson. Med. 64 (2) (2010) 457–471.

[12] Y. Yang, J. Sun, H. Li, Z. Xu, ADMM-CSNet: A deep learning approach for image compressive sensing, IEEE Trans. Pattern Anal. Mach. Intell. 42 (3) (2020) 521–538.

[13] J. Schlemper, J. Caballero, J.V. Hajnal, A. Price, D. Rueckert, A deep cascade of convolutional neural networks for MR image reconstruction, in: International Conference on Information Processing in Medical Imaging, 2017, pp. 647–658.

[14] K. Hammernik, T. Klatzer, E. Kobler, M.P. Recht, D.K. Sodickson, T. Pock, F. Knoll, Learning a variational network for reconstruction of accelerated MRI data, Magn. Reson. Med. 79 (6) (2017) 3055–3071.

[15] M. Akçakaya, S. Moeller, S. Weingärtner, K. Uğurbil, Scan-specific robust artificial-neural-networks for k-space interpolation (RAKI) reconstruction: Database-free deep learning for fast imaging, Magn. Reson. Med. 81 (1) (2019) 439–453.

[16] T.H. Kim, P. Garg, J.P. Haldar, LORAKI: Autocalibrated recurrent neural networks for autoregressive MRI reconstruction in k-space, 2019, arXiv preprint arXiv:1904.09390.

[17] Y. Arefeen, O. Beker, J. Cho, H. Yu, E. Adalsteinsson, B. Bilgic, Scan-specific artifact reduction in k-space (SPARK) neural networks synergize with physics-based reconstruction to accelerate MRI, Magn. Reson. Med. 87 (2) (2022) 764–780.

[18] S.A.H. Hosseini, C. Zhang, S. Weingärtner, S. Moeller, M. Stuber, K. Ugurbil, M. Akçakaya, Accelerated coronary MRI with sRAKI: A database-free self-consistent neural network k-space reconstruction for arbitrary undersampling, PLoS One 15 (2) (2020) 1–13.

[19] L. Zhang, X. Li, W. Chen, CAMP-Net: Context-aware multi-prior network for accelerated MRI reconstruction, 2023, arXiv preprint arXiv:2306.11238.

[20] H. Tao, W. Zhang, H. Wang, S. Wang, D. Liang, X. Xu, Q. Liu, Multi-weight respecification of scan-specific learning for parallel imaging, Magn. Reson. Imaging 97 (2023) 1–12.

[21] A.P. Leynes, N. Deveshwar, S.S. Nagarajan, P.E. Larson, Scan-specific self-supervised Bayesian deep non-linear inversion for undersampled MRI reconstruction, 2022, arXiv preprint arXiv:2203.00361.

[22] Y. Korkmaz, M. Yurt, S.U.H. Dar, M. Özbey, T. Cukur, Deep MRI reconstruction with generative vision transformers, in: International Workshop on Machine Learning for Medical Image Reconstruction, Springer, Cham, 2021, pp. 54–64.

[23] S. Arora, V. Roeloffs, M. Lustig, Untrained modified deep decoder for joint denoising and parallel imaging reconstruction, in: Proceedings of the 28th Annual Meeting of the ISMRM, 2020, p. 3585.

[24] M.Z. Darestani, R. Heckel, Accelerated MRI with un-trained neural networks, IEEE Trans. Comput. Imaging 7 (2021) 724–733.

[25] Y. Sui, O. Afacan, C. Jaimes, A. Gholipour, S.K. Warfield, Scan-specific generative neural network for MRI super-resolution reconstruction, IEEE Trans. Med. Imaging 41 (6) (2022) 1383–1399.

[26] R. Feng, Q. Wu, Y. Zhang, H. Wei, A scan-specific unsupervised method for parallel MRI reconstruction via implicit neural representation, 2022, arXiv preprint arXiv:2210.10439.

[27] D. Narnhofer, K. Hammernik, F. Knoll, T. Pock, Inverse GANs for accelerated MRI reconstruction, in: Proceedings of the SPIE, Vol. 11138, 2019, pp. 381–392.

[28] K.C. Tezcan, C.F. Baumgartner, R. Luechinger, K.P. Pruessmann, E. Konukoglu, MR image reconstruction using deep density priors, IEEE Trans. Med. Imaging 38 (7) (2019) 1633–1642.

[29] Q. Liu, Q. Yang, H. Cheng, S. Wang, M. Zhang, D. Liang, Highly undersampled magnetic resonance imaging using autoencoding priors, Magn. Reson. Med. 83 (1) (2020) 322–336.

[30] T. Eo, Y. Jun, T. Kim, J. Jang, H.-J. Lee, D. Hwang, KIKI-net: cross-domain convolutional neural networks for reconstructing undersampled magnetic resonance images, Magn. Reson. Med. 80 (5) (2018) 2188–2201.

[31] M. Mardani, E. Gong, J.Y. Cheng, S. Vasanawala, G. Zaharchuk, L. Xing, J.M. Pauly, Deep generative adversarial neural networks for compressive sensing MRI, IEEE Trans. Med. Imaging 38 (1) (2019) 167–179.

[32] H.K. Aggarwal, M.P. Mani, M. Jacob, MoDL: Model-Based deep learning architecture for inverse problems, IEEE Trans. Med. Imaging 38 (2) (2019) 394–405.

[33] S.U.H. Dar, M. Özbey, A.B. Çatlı, T. Çukur, A transfer-learning approach for accelerated MRI using deep neural networks, Magn. Reson. Med. 84 (2) (2020) 663–685.

[34] D. Lee, J. Yoo, S. Tak, J.C. Ye, Deep residual learning for accelerated MRI using magnitude and phase networks, IEEE Trans. Biomed. Eng. 65 (9) (2018) 1985–1995.

[35] P. Guo, J.M.J. Valanarasu, P. Wang, J. Zhou, S. Jiang, V.M. Patel, Over-and-under complete convolutional RNN for MRI reconstruction, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2021, pp. 13–23.

[36] G. Yiasemis, J.-J. Sonke, C. Sánchez, J. Teuwen, Recurrent variational network: A deep learning inverse problem solver applied to the task of accelerated MRI reconstruction, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 732–741.

[37] R. Hou, F. Li, IDPCNN: Iterative denoising and projecting CNN for MRI reconstruction, J. Comput. Appl. Math. 406 (2022) 113973.

[38] Z. Ramzi, G. Chaithya, J.-L. Starck, P. Ciuciu, NC-PDNet: a density-compensated unrolled network for 2D and 3D non-cartesian MRI reconstruction, IEEE Trans. Med. Imaging (2022).

[39] K. Kwon, D. Kim, H. Park, A parallel MR imaging method using multilayer perceptron, Med. Phys. 44 (12) (2017) 6209–6224.

[40] S. Wang, Z. Su, L. Ying, X. Peng, S. Zhu, F. Liang, D. Feng, D. Liang, Accelerating magnetic resonance imaging via deep learning, ISBI, 2016, pp. 514–517.

[41] J.C. Ye, Y. Han, E. Cha, Deep convolutional framelets: A general deep learning framework for inverse problems, SIAM J. Imaging Sci. 11 (2) (2018) 991–1048.

[42] J. Yoon, E. Gong, I. Chatnuntawech, B. Bilgic, J. Lee, W. Jung, J. Ko, H. Jung, K. Setsompop, G. Zaharchuk, E.Y. Kim, J. Pauly, J. Lee, Quantitative susceptibility mapping using deep neural network: Qsmnet, NeuroImage 179 (2018) 199–206.

[43] C.M. Hyun, H.P. Kim, S.M. Lee, S. Lee, J.K. Seo, Deep learning for undersampled MRI reconstruction, Phys. Med. Biol. 63 (13) (2018) 135007.

[44] A. Hauptmann, S. Arridge, F. Lucka, V. Muthurangu, J.A. Steeden, Real-time cardiovascular MR with spatio-temporal artifact suppression using deep learning–proof of concept in congenital heart disease, Magn. Reson. Med. 81 (2) (2019) 1143–1156.

[45] S.A.H. Hosseini, B. Yaman, S. Moeller, M. Hong, M. Akçakaya, Dense recurrent neural networks for accelerated MRI: History-cognizant unrolling of optimization algorithms, IEEE J. Sel. Top. Sign. Proces. 14 (6) (2020) 1280–1291.

[46] C. Qin, J. Schlemper, J. Caballero, A.N. Price, J.V. Hajnal, D. Rueckert, Convolutional recurrent neural networks for dynamic MR image reconstruction, IEEE Trans. Med. Imaging 38 (1) (2019) 280–290.

[47] T.M. Quan, T. Nguyen-Duc, W.-K. Jeong, Compressed sensing MRI reconstruction with cyclic loss in generative adversarial networks, IEEE Trans. Med. Imaging 37 (6) (2018) 1488–1497.

[48] S.U. Dar, M. Yurt, M. Shahdloo, M.E. Ildız, B. Tınaz, T. Çukur, Prior-guided image reconstruction for accelerated multi-contrast MRI via generative adversarial networks, IEEE J. Sel. Top. Sign. Proces. 14 (6) (2020) 1072–1087.

[49] Y. Chen, D. Firmin, G. Yang, Wavelet improved GAN for MRI reconstruction, in: Proceedings of SPIE, Medical Imaging 2021: Physics of Medical Imaging, Vol. 11595, 2021, 1159513.

[50] G. Elmas, S.U. Dar, Y. Korkmaz, E. Ceyani, B. Susam, M. Özbey, S. Avestimehr, T. Çukur, Federated learning of generative image priors for MRI reconstruction, IEEE Trans. Med. Imaging 42 (7) (2023) 1996–2009.

[51] M. Yaqub, F. Jinchao, S. Ahmed, K. Arshid, M.A. Bilal, M.P. Akhter, M.S. Zia, GAN-TL: Generative adversarial networks with transfer learning for mri reconstruction, Appl. Sci. 12 (17) (2022) 8841.

[52] P. Guo, Y. Mei, J. Zhou, S. Jiang, V.M. Patel, Reconformer: Accelerated mri reconstruction using recurrent transformer, 2022, arXiv preprint arXiv:2201.09376.

[53] A. Güngör, S.U. Dar, Ş. Öztürk, Y. Korkmaz, G. Elmas, M. Özbey, T. Çukur, Adaptive diffusion priors for accelerated MRI reconstruction, Med. Image Anal. 88 (2023) 102872.

[54] C. Peng, P. Guo, S.K. Zhou, V.M. Patel, R. Chellappa, Towards performant and reliable undersampled MR reconstruction via diffusion model sampling, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer, 2022, pp. 623–633.

[55] K. Wang, J.I. Tamir, A. De Goyeneche, U. Wollner, R. Brada, S.X. Yu, M. Lustig, High fidelity deep learning-based MRI reconstruction with instance-wise discriminative feature matching loss, Magn. Reson. Med. 88 (1) (2022) 476–491.

[56] F. Gadjimuradov, T. Benkert, M.D. Nickel, A. Maier, Robust partial Fourier reconstruction for diffusion-weighted imaging using a recurrent convolutional neural network, Magn. Reson. Med. 87 (4) (2022) 2018–2033.

[57] Z. Ramzi, C. G R, J.-L. Starck, P. Ciuciu, NC-PDNet: A density-compensated unrolled network for 2D and 3D non-cartesian MRI reconstruction, IEEE Trans. Med. Imaging 41 (7) (2022) 1625–1638.

[58] D. Polak, S. Cauley, B. Bilgic, E. Gong, P. Bachert, E. Adalsteinsson, K. Setsompop, Joint multi-contrast variational network reconstruction (jVN) with application to rapid 2D and 3D imaging, Magn. Reson. Med. 84 (3) (2020) 1456–1469.

[59] J.Y. Cheng, M. Mardani, M.T. Alley, J.M. Pauly, S.S. Vasanawala, DeepSPIRiT: Generalized parallel imaging using deep convolutional neural networks, in: Proceedings of the 26th Annual Meeting of the ISMRM, 2018, p. 0570.

[60] K. Pooja, Z. Ramzi, G. Chaithya, P. Ciuciu, MC-PDNet: Deep unrolled neural network for multi-contrast Mr image reconstruction from undersampled K-space data, in: 2022 IEEE 19th International Symposium on Biomedical Imaging, ISBI, 2022, pp. 1–5.

[61] Z. Fabian, B. Tinaz, M. Soltanolkotabi, Humus-net: Hybrid unrolled multi-scale network architecture for accelerated mri reconstruction, Adv. Neural Inf. Process. Syst. 35 (2022) 25306–25319.

[62] I.P. Machado, E. Puyol-Antón, K. Hammernik, G. Cruz, D. Ugurlu, I. Olakorede, I. Oksuz, B. Ruijsink, M. Castelo-Branco, A.A. Young, et al., A deep learning-based integrated framework for quality-aware undersampled cine cardiac MRI reconstruction and analysis, 2022, arXiv preprint arXiv:2205.01673.

[63] A. Aghabiglou, E.M. Eksioglu, Deep unfolding architecture for MRI reconstruction enhanced by adaptive noise maps, Biomed. Signal Process. Control 78 (2022) 104016.

[64] F. Knoll, K. Hammernik, E. Kobler, T. Pock, M.P. Recht, D.K. Sodickson, Assessment of the generalization of learned image reconstruction and the potential for transfer learning, Magn. Reson. Med. 81 (1) (2019) 116–128.

[65] A.S. Chaudhari, C.M. Sandino, E.K. Cole, D.B. Larson, G.E. Gold, S.S. Vasanawala, M.P. Lungren, B.A. Hargreaves, C.P. Langlotz, Prospective deployment of deep learning in MRI: A framework for important considerations, challenges, and recommendations for best practices, J. Magn. Reson. Imaging 54 (2) (2021) 357–371.

[66] N. Tavaf, A. Torfi, K. Ugurbil, P.-F. Van de Moortele, GRAPPA-GANs for parallel MRI reconstruction, 2021, arXiv preprint arXiv:2101.03135.

[67] C.M. Sandino, P. Lai, S.S. Vasanawala, J.Y. Cheng, Accelerating cardiac cine MRI using a deep learning-based ESPIRiT reconstruction, Magn. Reson. Med. 85 (1) (2021) 152–167.

[68] S.U.H. Dar, M. Yurt, T. Çukur, A few-shot learning approach for accelerated MRI via fusion of data-driven and subject-driven priors, in: Proceedings of the 29th Annual Meeting of the ISMRM, 2021, p. 1949.

[69] J.I. Tamir, S.X. Yu, M. Lustig, Unsupervised deep basis pursuit: Learning reconstruction without ground-truth data, in: Proceedings of the 27th Annual Meeting of the ISMRM, 2019, p. 0660.

[70] E.K. Cole, F. Ong, S.S. Vasanawala, J.M. Pauly, Fast unsupervised MRI reconstruction without fully-sampled ground truth data using generative adversarial networks, in: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops, 2021, pp. 3988–3997.

[71] B. Yaman, S.A.H. Hosseini, S. Moeller, J. Ellermann, K. Uğurbil, M. Akçakaya, Self-supervised learning of physics-guided reconstruction neural networks without fully sampled reference data, Magn. Reson. Mater. 84 (6) (2020) 3172–3191.

[72] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, U.S. Kamilov, RARE: Image reconstruction using deep priors learned without groundtruth, IEEE J. Sel. Top. Sign. Proces. 14 (6) (2020) 1088–1099.

[73] S. Wang, R. Wu, C. Li, J. Zou, Z. Zhang, Q. Liu, Y. Xi, H. Zheng, PARCEL: Physics-based unsupervised contrastive representation learning for multi-coil MR imaging, 2022, arXiv:2202.01494.

[74] S.A.H. Hosseini, C. Zhang, K. Uğurbil, S. Moeller, M. Akçakaya, sRAKI-RNN: accelerated MRI with scan-specific recurrent neural networks using densely connected blocks, in: Wavelets and Sparsity XVIII, Vol. 11138, International Society for Optics and Photonics, 2019, p. 111381B.

[75] Q. Huang, Y. Xian, D. Yang, H. Qu, J. Yi, P. Wu, D.N. Metaxas, Dynamic MRI reconstruction with end-to-end motion-guided network, Med. Image Anal. 68 (2021) 101901.

[76] J. Huang, S. Wang, G. Zhou, W. Hu, G. Yu, Evaluation on the generalization of a learned convolutional neural network for MRI reconstruction, Magn. Reson. Imaging 87 (2022) 38–46.

[77] M. Uecker, P. Lai, M.J. Murphy, P. Virtue, M. Elad, J.M. Pauly, S.S. Vasanawala, M. Lustig, ESPIRiT-an eigenvalue approach to autocalibrating parallel MRI: Where SENSE meets GRAPPA, Magn. Reson. Med. 71 (3) (2014) 990–1001.

[78] D. Ulyanov, A. Vedaldi, V. Lempitsky, Deep image prior, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2018, pp. 9446–9454.

[79] F. Knoll, J. Zbontar, A. Sriram, M.J. Muckley, M. Bruno, A. Defazio, M. Parente, K.J. Geras, J. Katsnelson, H. Chandarana, Z. Zhang, M. Drozdzalv, A. Romero, M. Rabbat, P. Vincent, J. Pinkerton, D. Wang, N. Yakubova, E. Owens, C.L. Zitnick, M.P. Recht, D.K. Sodickson, Y.W. Lui, fastMRI: A publicly available raw k-space and DICOM dataset of knee images for accelerated MR image reconstruction using machine learning, Radiol. Artif. Intell. 2 (1) (2020) e190007.

[80] T. Zhang, J.M. Pauly, S.S. Vasanawala, M. Lustig, Coil compression for accelerated imaging with Cartesian sampling, Magn. Reson. Med. 69 (2) (2013) 571–582.

[81] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, in: Advances in Neural Information Processing Systems, Vol. 33, 2020, pp. 6840–6851.

[82] D.P. Kingma, J.L. Ba, Adam: a method for stochastic optimization, in: International Conference on Learning Representations, 2015.

[83] O. Dalmaz, U. Mirza, G. Elmas, M. Özbey, S.U. Dar, E. Ceyani, S. Avestimehr, T. Çukur, One model to unite them all: Personalized federated learning of multi-contrast MRI synthesis, 2022, arXiv:2207.06509.

[84] S.A.H. Hosseini, B. Yaman, S. Moeller, M. Akçakaya, High-fidelity accelerated MRI reconstruction by scan-specific fine-tuning of physics-based neural networks, in: 2020 42nd Annual International Conference of the IEEE Engineering in Medicine Biology Society, EMBC, 2020, pp. 1481–1484.

[85] K. Lei, M. Mardani, J.M. Pauly, S.S. Vasanawala, Wasserstein GANs for MR imaging: from paired to unpaired training, IEEE Trans. Med. Imaging 40 (1) (2021) 105–115.

[86] C. Zhang, S.A.H. Hosseini, S. Moeller, S. Weingärtner, K. Ugurbil, M. Akçakaya, Scan-specific residual convolutional neural networks for fast MRI using residual RAKI, in: 2019 53rd Asilomar Conference on Signals, Systems, and Computers, 2019, pp. 1476–1480.

[87] M. Murphy, M. Alley, J. Demmel, K. Keutzer, S. Vasanawala, M. Lustig, Fast l_1-SPIRiT compressed sensing parallel imaging MRI: scalable parallel implementation and clinically feasible runtime, IEEE Trans. Med. Imaging 31 (6) (2012) 1250–1262.

[88] M. Özbey, O. Dalmaz, S.U. Dar, H.A. Bedel, Ş. Özturk, A. Güngör, T. Çukur, Unsupervised medical image translation with adversarial diffusion models, IEEE Trans. Med. Imaging (2023).

[89] O. Dalmaz, M. Yurt, T. Çukur, ResViT: Residual vision transformers for multi-modal medical image synthesis, IEEE Trans. Med. Imaging 41 (7) (2022) 2598–2614.

[90] M.U. Mirza, O. Dalmaz, H.A. Bedel, G. Elmas, Y. Korkmaz, A. Gungor, S.U. Dar, T. Çukur, Learning Fourier-constrained diffusion bridges for MRI reconstruction, 2023.

[91] H. Zhou, J. Hou, Y. Zhang, J. Ma, H. Ling, Unified gradient- and intensity-discriminator generative adversarial network for image fusion, Inf. Fusion 88 (2022) 184–201.

[92] J. Liu, R. Dian, S. Li, H. Liu, SGFusion: A saliency guided deep-learning framework for pixel-level image fusion, Inf. Fusion 91 (2023) 205–214.