Appl Magn Reson (2026) 57:41
https://doi.org/10.1007/s00723-026-01862-8

Applied
Magnetic Resonance

RESEARCH

SPIRiT Regularization: Parallel MRI with a Combination of Sensitivity Encoding and Linear Predictability

Nicholas Dwork[1] · **Alex McManus**[2] · **Stephen Becker**[2] · **Gennifer T. Smith**[3]

Received: 28 January 2026 / Revised: 19 May 2026 / Accepted: 10 July 2026
© The Author(s), under exclusive licence to Springer-Verlag GmbH Austria, part of Springer Nature 2026

Abstract

This manuscript presents a novel combination of technologies that yield images of improved quality for accelerated Magnetic Resonance Imaging (MRI). Two established methods for accelerating MRI include *parallel imaging* and *compressed sensing*. Two types of parallel imaging include *linear predictability*, which assumes that the Fourier samples are linearly related, and *sensitivity encoding*, which incorporates *a priori* knowledge of the sensitivity maps. In this work, we combine compressed sensing with both types of parallel imaging using a novel regularization term: *SPIRiT regularization*. For a given number of samples, the images reconstructed with SPIRiT regularization are improved over those reconstructed without it. We demonstrate these improvements on data of a knee, a brain, and an ankle.

1 Introduction

Magnetic Resonance Imaging (MRI) is a Fourier sensing modality, meaning the data are collected in the spatial frequency domain. MRI images without any ionizing radiation and presents excellent contrast between soft tissues. However, compared to other imaging methods, such as x-ray computed tomography (CT) or ultrasound, which offer comparable resolution, MRI is slow to collect data. Whereas an unaccelerated 3D MRI scan for a volume $25 \times 25 \times 25$ cm^3 will require >5 min, the same volume can be imaged with CT in <10 s. Its slow scan time makes the cost of each imaging study high and may necessitate the use of anesthesia for those patients who have difficulty remaining still (e.g., pediatric subjects or subjects in pain). One approach to accelerating MRI is to reduce the number of data samples collected.

✉ Nicholas Dwork
nicholas.dwork@cuanschutz.edu

[1] Department of Biomedical Informatics, University of Colorado School of Medicine, Aurora 80045, USA

[2] Department of Applied Mathematics, University of Colorado Boulder, Boulder 80309, USA

[3] Department of Engineering, University of San Francisco, San Francisco 94117, USA

Published online: 19 July 2026

The accompanying challenge is to reconstruct a high-quality image from those fewer data collected. Established technologies for accelerating MRI in this way include partial Fourier sampling [1–4], parallel imaging [5–7], compressed sensing [8–10], and incorporating knowledge of the object's support [11–14]. Parallel imaging, which takes advantage of multiple sensing coils with distinct sensitivity profiles, is now a mainstay clinical method for reducing scan time.

There are two main lines of research on parallel MRI: *linear predictability* and *sensitivity encoding*. Parallel Imaging with Linear Predictability (PILP) assumes that each Fourier value is equal to a linear combination of nearby Fourier values [15–18]. Sensitivity encoding incorporates *a priori* knowledge of the coils' sensitivity maps into the reconstruction [6, 19–21]. Sensitivity encoding methods often reconstruct the image by numerically solving an optimization problem that requires coil sensitivities [22]. A common method for estimating sensitivity maps is ESPIRiT [23], which is equivalent to PISCO [24].

ESPIRiT is an extremely successful technique (available through the BART software package [25]) which was described as a method "where SENSE meets GRAPPA" [23]; SENSE was an early sensitivity encoding method [19] that reconstructed the image by applying the pseudoinverse of the system matrix, and GRAPPA was an early PILP method [16]. This description is fitting: ESPIRiT leverages the assumptions of PILP to estimate sensitivity maps, which are then used to reconstruct the image through a sensitivity encoding approach. The estimated sensitivity maps implicitly encode linear predictability.

The acceleration rates achievable for parallel MRI with pseudoinverse-based methods are limited by the condition of the system matrix. Higher acceleration requires incorporating additional prior information into the reconstruction. Compressed sensing employs an ℓ_1 penalty on sparsifying transform coefficients, permitting higher acceleration while still reconstructing high quality images. While compressed sensing is has been combined with linear predictability and sensitivity encoding separately, it has yet to be combined with sensitivity encoding and linear predictability into a single reconstruction method.

Since sensitivity encoding reconstructions enforce data consistency (i.e., agreement with measured Fourier samples), the reconstruction implicitly enforces SPIRiT self-consistency at the sampled k-space locations [18]. However, data consistency alone does not constrain SPIRiT self-consistency at unsampled frequencies, where the solution is determined by the choice of regularization and the effective nullspace of the forward model. In this manuscript, we make this distinction explicit and introduce SPIRiT regularization, which enforces SPIRiT self-consistency across the entirety of k-space, complementing the implicit enforcement that arises from data consistency. We demonstrate that adding this regularization into a parallel imaging and compressed sensing (PICS) reconstruction system improves image quality compared to PICS alone.

2 Methods

For this work, we are assuming that data is collected on a Cartesian grid with two dimensions of phase encodes and one dimension of readout, and with all readout lines parallel to each other. (This assumption is not necessary; future work will investigate non-Cartesian 3D data collection trajectories.) The data is fully sampled along the readout direction and undersampled in the phase encoded dimensions. After data collection, an inverse Fourier transform is applied along each readout line to place the data into the hybrid (k_x, k_y, z) domain [26]. Once in the hybrid domain, each slice can be reconstructed independently; that is the approach we will take here.

SPIRiT [27] makes the fundamental assumption of PILP: any Fourier value equals a linear combination of nearby Fourier values with linear coefficients that are invariant across the frequency domain. In [17], Haldar et al. provide a necessary and sufficient condition for when PILP is accurate. While valuable, it is difficult to relate this condition to the physical system. In [18] it was shown that PILP is accurate when there exists a linear combination of the coils' sensitivities such that the result is a complex exponential. More specifically, consider the set of $c = 1, 2, \ldots, C$ coils. PILP accurately interpolates the Fourier value for coil ℓ at spatial frequency $k = (k_x, k_y)$ whenever there exists a $d = (d_x, d_y)$ where the Fourier value at $k + d$ is known and there exist coefficients $n(d) \in \mathbb{C}^C$ such that

$$\sum_{c=1}^{C} n^{(c)}(d)\, S^{(c)}(x, y) \approx e^{i\left(d_x x + d_y y\right)} S^{(\ell)}(x, y), \tag{1}$$

where $S^{(\ell)}$ is the sensitivity of the ℓ^{th} coil and $n^{(\ell)}(0) = 0$.

SPIRiT uses a fully-sampled auto-calibration region (ACR) to estimate the optimal linear coefficients $\mathcal{W}$ for a set of directions that lie within a kernel. Since the coefficients do not vary across the frequency domain, the process of interpolation can be expressed as sum of convolution operations:

$$FS^{(\ell)}m = \sum_{c=1}^{C} \mathcal{W}_{(\ell)}^{(c)} \circledast FS^{(c)}m, \tag{2}$$

where $m \in \mathbb{C}^{U \times V}$ represents the image, F represents the Discrete Fourier Transform ([DFT] – as defined in App. A), $\circledast$ represents circular convolution, and $\mathcal{W}_{(\ell)}^{(c)}$ is the kernel for interpolating the ℓ^{th} coil's Fourier values. Since the sum of convolutions is a linear transformation, Eq. (2) can be expressed as

$$\left(W^{(\ell)} - I\right) F\, S^{(\ell)}\, m = 0, \tag{3}$$

where $W^{(\ell)}\, k = \sum_{c=1}^{C} \mathcal{W}_{(\ell)}^{(c)} \circledast k$. The values of $W^{(\ell)}$ are determined by solving the following least-squares optimization problem:

$$\underset{W^{(\ell)}}{\text{minimize}} \left\| W^{(\ell)} * k_{\text{ACR}} - k_{\text{interior(ACR)}} \right\|_2. \tag{4}$$

Here, $*$ represents convolution and k_{ACR} is a vector of the Fourier values of the spatial frequencies located within the ACR. Problem Eq. (4) is a least-squares problem that can be solved with LSQR, which avoids explicitly forming the large matrix. Note that we are using convolution instead of circular convolution and we are eliminating any values from the comparison where the kernel $W^{(\ell)}$ extends outside of the ACR by restricting the objective to those values on the interior of the ACR. See [18] for further details.

2.1 Problem Setup

An image can be reconstructed with parallel imaging and compressed sensing by solving the Lagrangian form of a basis pursuit denoising optimization problem [28]:

$$\underset{m}{\text{minimize}} \ (1/2)\|\boldsymbol{M}\boldsymbol{F}\boldsymbol{S}\,m - \boldsymbol{b}\|_2^2 \ + \nu\|\Psi m\|_1. \tag{5}$$

Here, $\|\cdot\|_p$ is the ℓ_p norm, $\boldsymbol{S} = \left(S^{(1)}, S^{(2)}, \cdots, S^{(C)}\right)$ is a block-column matrix where each block element represents multiplication by a coil's sensitivity map, $\boldsymbol{F} = \text{diag}(F, F, \cdots, F)$ is a block-diagonal matrix that performs the DFT on each modulated image,

$\boldsymbol{M} = \text{diag}(M, M, \cdots, M)$ is a block-diagonal matrix where M is the data sampling mask that isolates the values for those frequencies that were collected, $\boldsymbol{b} = \left(b^{(1)}, b^{(2)}, \cdots b^{(c)}\right)$ is a complex block-column vector where $b^{(c)}$ is the vector of data collected by the c^{th} coil, $\nu > 0$ is the compressed sensing regularization parameter, and Ψ is a sparsifying transformation. The optimization variable m is a complex vectorized image (i.e., the columns of the image are concatenated into a single vector). For this work, we use the Daubechies-4 discrete wavelet transformation as the sparsifying transformation [29].

Equation 3 is of the form that could be included as a regularization function in an optimization problem. The motivation for adding the SPIRiT term is that PICS and SPIRiT encode complementary prior information. The PICS data-consistency term enforces agreement with the measured samples, while the wavelet penalty promotes sparsity of the reconstructed image in the transformed domain. However, neither term explicitly enforces the multi-coil k-space self-consistency relations learned from the autocalibration region. This leaves the missing k-space values determined only indirectly by the wavelet sparsity prior. SPIRiT regularization addresses this specific gap. The SPIRiT kernels estimate coil-to-coil linear predictability relations from the fully sampled autocalibration region. Adding the regularization term in enforces these learned relations throughout k-space, including at unsampled locations.

Note that Eq. (3) is true in the ideal case where Eq. (1) is perfectly satisfied; i.e., where the linear combination of sensitivity maps equals the ℓ^{th} sensitivity modulated by a perfect complex exponential. However, it is unlikely that this is exactly the case. Instead, the right-hand side of Eq. (3) will be nonzero, but small. With this understanding, parallel imaging with sensitivity encoding can be combined with PILP and compressed sensing as the following optimization problem:

$$\underset{m}{\text{minimize}} \ (1/2)\|\boldsymbol{M}\boldsymbol{F}\boldsymbol{S}\,m - \boldsymbol{b}\|_2^2 \ + \nu\|\Psi m\|_1 \ + \frac{\lambda_s}{2\kappa}\sum_{c=1}^{C}\left\|\left(W^{(c)} - I\right)F\,S^{(c)}\,m\right\|_{\gamma,2}^2. \tag{6}$$

The new regularization term in Eq. (6) is *SPIRiT regularization*. It provides a principled extension to PICS by explicitly enforcing multi-coil self-consistency, a structure inherent to parallel MRI that is not otherwise enforced by the standard formulation. Note that the data consistency term enforces agreement with measured samples only on the sampling mask M. When the sensitivity maps satisfy linear predictability (as when they are estimated by ESPIRiT or PISCO), this implies an implicit enforcement of SPIRiT self-consistency at sampled k-space locations. SPIRiT regularization (SR) explicitly enforces self-consistency at unsampled frequencies as well.

In problem (6), $\|\cdot\|_{\gamma,2}$ is the weighted ℓ_2 norm, and $\lambda_s \geq 0$ is the spirit regularization parameter.[1] The weighted ℓ_2 norm is defined as follows:

$$\|x\|_{\gamma,2} = \sqrt{\gamma_1|x_1|^2 + \gamma_2|x_2|^2 + \cdots + \gamma_N|x_N|^2}.$$

The value of κ is used to balance the terms of the objective function and bring the appropriate value for λ_s closer to 1; it is set to the square root of the ratio of matrix induced norms:

$$\kappa = \sqrt{\frac{\left\|\text{diag}\left(\left(W^{(1)} - I\right)F\,S^{(1)}, \left(W^{(2)} - I\right)F\,S^{(2)}, \ldots, \left(W^{(C)} - I\right)F\,S^{(C)}\right)\right\|}{\|\boldsymbol{M}\boldsymbol{F}\boldsymbol{S}\|}}.$$

In Eq. (6), without using a weighted norm (i.e., $\gamma = (1, 1, \ldots, 1)$), the optimization algorithm could make the errors for large frequencies as large as the errors for small frequencies to yield an optimal solution. Since the magnitudes of Fourier values of natural images decrease with increasing frequency, as a percentage of the true Fourier magnitudes, these errors would be a much larger percentage of the Fourier values in the higher frequencies than in the lower frequencies. To mitigate this outcome, we alter the weights of the weighted norm.

The magnitude spectrum of natural images (without noise) for frequencies other than 0 tends to follow a power law [33–35]. (The power law would predict infinite power for the 0 frequency, which is never valid for natural images.) We will weight the norm with the inverse of a fit based on the power law. To accommodate the presence of noise, we will independently fit small frequencies (those with magnitude greater than 0 and less than 0.1) and large frequencies (those with magnitude greater than 0.1) to a power law and then take the maximum of the fit, as follows:

$$P(k) = \max\left(m_L \cdot |k|^{-p_L}, m_H \cdot |k|^{-p_H}\right).$$

[1] Note that in this formulation, we are assuming that the samples collected are located on a Cartesian grid. If this were not the case, then MF would be replaced with a non-uniform DFT [30–32]

To identify the low and high scaling coefficients m_L and m_H, and the corresponding exponents p_L and p_H, we use the Levenberg-Marquardt optimization algorithm [36] to find the values that minimize the ℓ_2 norm of the difference between $P(k)$ and $|b(k)|$ for all collected spatial frequencies. To identify the value of the fit at $k = 0$, we fit a line to a small number of points surrounding the 0 frequency. Once the fit is complete, the values of γ are set to the inverse of the fitted values:

$$\gamma_i = P(k_i)^{-1} = \min\left(\frac{|k_i|^{p_L}}{m_L}, \frac{|k_i|^{p_H}}{m_H}\right). \tag{7}$$

This weighting ensures that the errors in the high frequencies are much smaller than those in the low frequencies, commensurate with the magnitude of the Fourier values. Reconstructing an MR image by solving problem (6) with weights determined according to Eq. (7) is parallel imaging with compressed sensing and SPIRiT regularization (PICS+SR).

2.2 Solving the Optimization Problems

The optimization problem of (6) can be solved with the Fast Iterative Shrinkage Thresholding Algorithm (FISTA) [37]. Specifically, we will use FISTA with line search as detailed in [38]. FISTA solves problems of the form

$$\text{minimize} \;\; G(y) + H(y), \tag{8}$$

where G is convex and differentiable, and H is proper, convex, and lower semicontinuous with a computationally feasible proximal operator, defined as follows:

$$\text{prox}_{tH}(x) = \underset{v}{\text{argmin}}\, H(v) + \frac{1}{2t}\|x - v\|_2^2.$$

As an example, consider the PICS problem of (5). Let $G(m) = (1/2)\|\boldsymbol{MFS}m - \boldsymbol{b}\|_2^2$ and let $H(m) = \nu\|\Psi m\|_1$. Then G is differentiable, H has a simple proximal operator when Ψ is a tight frame, and problem (8) is the PICS problem, which can now be solved with FISTA. The proximal operator of H is soft-thresholding applied to the wavelet coefficients of m [39]. We let Ψ be the discrete Daubechies-4 wavelet transform, which is orthogonal (and, therefore, tight frame), so H has a simple proximal operator.

If we instead define $G(m) = (1/2)\|\boldsymbol{MFS}m - \boldsymbol{b}\|_2^2 + (\lambda_s/2)\sum_{c=1}^{C}\left\|\left(W^{(c)} - I\right)F\,S^{(c)}\,m\right\|_{\gamma,2}^2$ with the same definition of H, then G is differentiable and Eq. (8) becomes the PICS+SR problem of (6), which can be solved with FISTA.

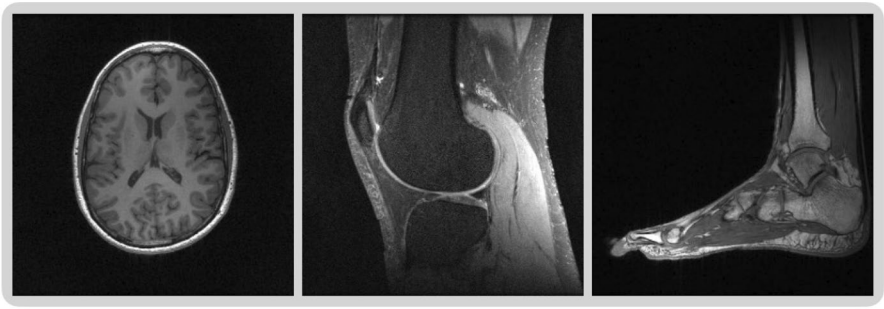

Fig. 1 Fully-sampled reconstructions of the brain, knee, and ankle images analyzed in this work

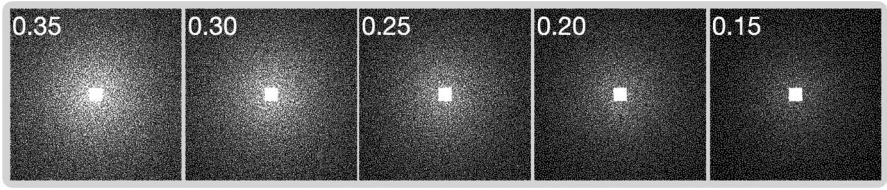

Fig. 2 The sampling patterns used in this work ranging from sample fractions of 0.15 to 0.35 where a white spot indicates a sample that was collected and a black spot indicates a sample that was not. The white numbers indicate the fraction of all samples that were collected. The white square in the middle is the auto-calibration region (ACR)

3 Experiments

We demonstrate the effectiveness of the proposed algorithm on MRI data of a knee (made public through mridata.org [40]), and a brain and an ankle (made public with [41]); all images were size 256×256. For these results, each dataset was scaled so that its maximum magnitude was 1. The sensitivity maps were identified with PISCO [24]. Figure 1 shows the fully-sampled reconstructions.

All data were collected with 3DFT trajectories; i.e., the samples were located on a Cartesian grid and there were two dimensions of phase encoding with one dimension of readout. To assess the quality that could be attained with accelerated MRI, the data was retrospectively downsampled according to a variable density Poisson disc sampling pattern that included a fully-sampled region centered on the 0 frequency – the autocalibration region (ACR) – generated according to [42]. The size of the ACR was determined according to the two-level sampling pattern of [43] and was used in [41, 44, 45]; with the four scales that we used for the wavelet transformation for this work, the size of the finest scale and the size of the ACR is 16×16. The sampling patterns for the different sample fractions studied in this manuscript are shown in Fig. 2.

To quantify image quality, we compare the reconstructed image to the fully-sampled reconstruction: the result of solving problem (5) with $\nu = 0$ without any

downsampling. We compute the Structural Similarity Index Measure (SSIM) [46] between the magnitude of the reconstructed image and the fully-sampled reconstruction, and we compute the peak signal-to-noise ratio (PSNR) between the two complex images [47].

4 Results

Images were reconstructed for all combinations of values of $\nu \in \{0.01, 0.1, 1, 10\}$ and $\lambda_s \in \{0.1, 0.5, 1, 2, 4, 5, 10\}$. Results for the reconstructions that yielded the highest value of SSIM are reported here. Quality metric values for all three data-cases—the brain, knee, and ankle—are shown in Table 1. For all results, the optimization algorithm was run for 1000 iterations; the difference in objective value between 500 and 1000 iterations was negligible (data not shown), indicating

Table 1 Quality metric values

Dataset	Metric	Sample fraction	PICS	PICS+SR
Brain	SSIM	0.35	0.966	0.967
		0.3	0.962	0.964
		0.25	0.959	0.959
		0.2	0.952	0.952
	PSNR (dB)	0.35	40.15	40.24
		0.3	39.60	39.69
		0.25	39.06	39.15
		0.2	38.38	38.47
Knee	SSIM	0.35	0.800	0.814
		0.3	0.779	0.796
		0.25	0.750	0.772
		0.2	0.721	0.747
	PSNR (dB)	0.35	29.67	30.42
		0.3	29.12	29.93
		0.25	28.49	29.30
		0.2	27.89	28.67
Ankle	SSIM	0.35	0.974	0.975
		0.3	0.970	0.971
		0.25	0.966	0.967
		0.2	0.960	0.961
	PSNR (dB)	0.35	40.74	40.83
		0.3	39.94	40.03
		0.25	39.06	39.16
		0.2	37.95	38.06

Quality metric values for reconstructions using parallel imaging and compressed sensing (PICS) and PICS with SPIRiT regularization (PICS+SR). The best reconstruction with respect to each metric is underlined

convergence. Incorporating our novel SPIRiT regularization consistently improved both SSIM and PSNR .

Figure 3 shows reconstructions of the ankle with and without SPIRiT regularization (SR) for several different sampling fractions. As expected, the image quality of PICS and PICS+SR improves with sample fraction. For all sample fractions, the inclusion of SPIRiT regularization improves the image quality. The improvement in image quality when using SPIRiT regularization is subtle and presents as a general reduction in the noise across the image (rather than an improvement in any localized structural features). The improvement with SPIRiT regularization is more significant at lower sample fractions.

Figure 4 shows the reconstructions of the knee with a 0.30 sample fraction with the magnitudes of the differences multiplied by 8. The difference images show the improvement gained by including SPIRiT regularization in the reconstruction.

Figure 5 shows the reconstructions of the brain with a 0.20 sample fraction along with the magnitudes of the differences multiplied by 8. The improvement in the brain with SPIRiT regularization is more subtle than that of the knee, as can be seen in the difference images.

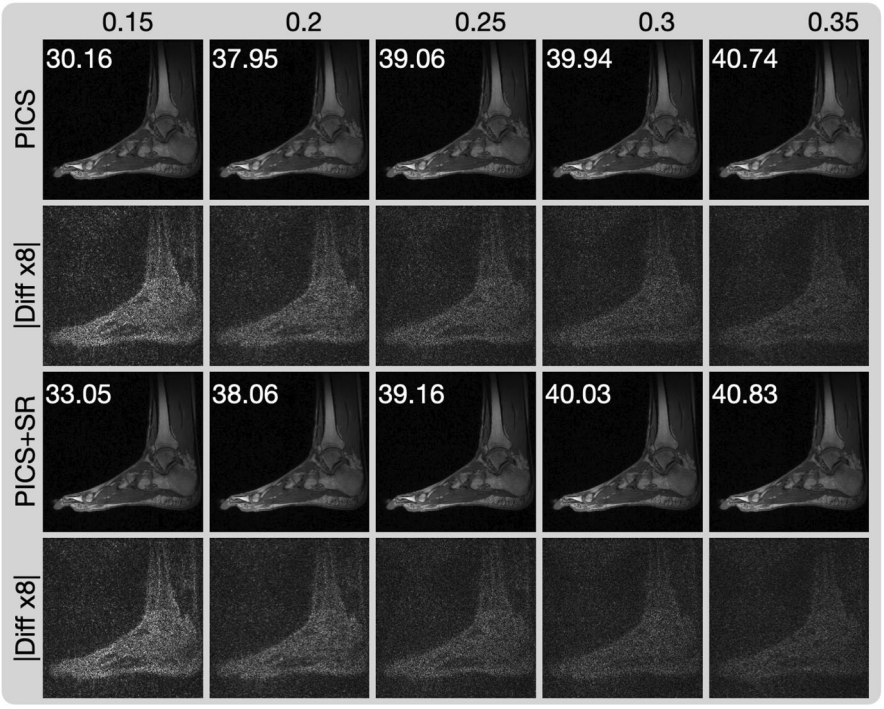

Fig. 3 Reconstructions of the ankle for several different sampling fractions. The first and third rows show the reconstructions with parallel imaging and compressed sensing (PICS), and PICS with SPIRiT regularization, respectively. The second and fourth rows show the magnitude of the difference images multiplied by 8. The PSNR in decibels for each reconstruction is also presented in white

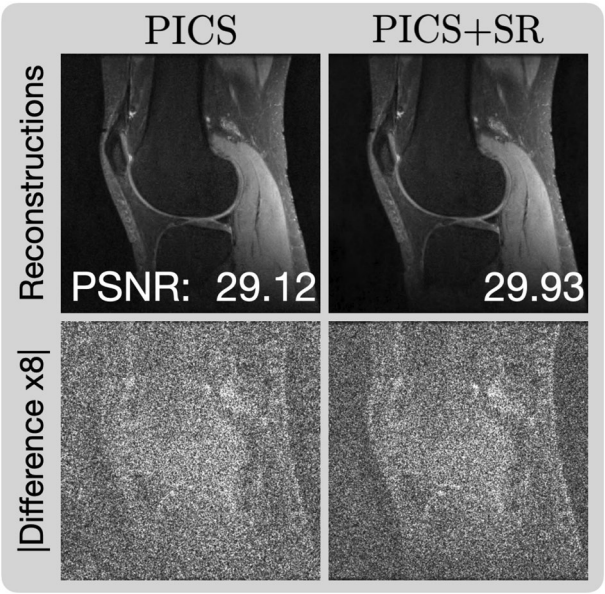

Fig. 4 Reconstructions of the knee with a 0.3 sample fraction using parallel imaging and compressed sensing (PICS) and PICS with SPIRiT regularization (SR). The second row shows the magnitudes of the differences between the under-sampled reconstructions and the fully-sampled reconstructions multiplied by 8. The PSNR in decibels for each reconstruction is also presented in white

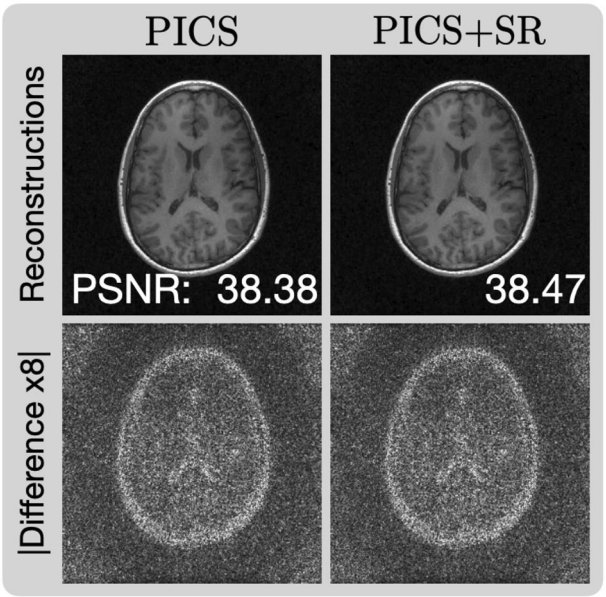

Fig. 5 Reconstructions of the brain with a 0.2 sample fraction using parallel imaging and compressed sensing (PICS), PICS with SPIRiT regularization (SR). The second row shows the magnitudes of the differences between the under-sampled reconstructions and the fully-sampled reconstructions multiplied by 8. The PSNR in decibels for each reconstruction is also presented in white

5 Conclusion

When combined with PICS, SPIRiT regularization improves image quality over the use of PICS alone. The improvement is more significant at lower sample fractions. In these experiments, sensitivity maps were estimated with PISCO (which is equivalent to ESPIRiT) and, therefore, incorporates linear predictability into the sensitivity model. As a result, the data consistency term already implicitly enforces SPIRiT self-consistency at sampled k-space locations. SPIRiT regularization complements this by explicitly enforcing self-consistency at unsampled frequencies, yielding a subtle, but consistent improvement. Importantly, the only change comes in the form of the reconstruction algorithm; the process of data collection remains the same. Thus, this improvement comes without any additional scan time or modifications to the clinical workflow.

In all cases, for both PICS and PICS+SR, the λ and γ values that yielded the optimal PSNR values were 0.1 and 0.5. This suggests that we could use the same values for new datasets and expect a high quality reconstruction. However, a study of much more data would be required to make such a conclusion.

For this work, we investigated an additional modification where knowledge of the support was included in the reconstruction process; our approach is detailed in App. B. The support is the set of pixels that image anatomy. While inclusion of the support did not improve image quality significantly in this work, in past work, it was shown that the sampling pattern could be reduced using the support alone [14]. Thus, it may be more beneficial to incorporate knowledge of the support with something other than a variable density Poisson disc sampling pattern; perhaps a variable density subsampling of the pattern presented in [14] would be more appropriate. Separately, the method of [24] identifies a conservative superset of the support. In particular, for the knee, only a small portion of the region outside of the knee is excluded from the support (see Fig. 6 in App. B). A different method that identifies more of the region outside of the support (i.e., outside of the anatomy) may improve the effect of including the support in the reconstruction method. We will explore these ideas in future work.

Based on the previous understanding provided in [18], SPIRiT regularization is appropriate when the sensitivity maps linearly combine into complex exponentials. This is more likely to be the case when a larger number of smaller coils are used for sensing.

This work presented results with 3DFT trajectories, which allowed for separate reconstructions of each slice in the volume. However, there is nothing inherent in this method that requires 3DFT trajectories. Future work will analyze SPIRiT regularization in the context of non-Cartesian trajectories and/or dynamic imaging.

In this work, we fixed the number of iterations of the optimization algorithms to 1000. Rather than having a fixed number of iterations for each optimization algorithm, it may be beneficial to include a dynamic stopping criteria that would automatically stop the optimization algorithms when close to convergence. This could reduce the time to reconstruct the image without any reduction in image quality.

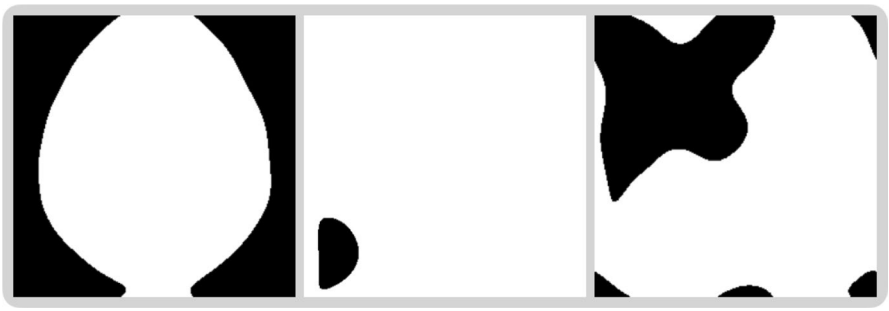

Fig. 6 The support regions identified with PISCO for the brain (left), the knee (center), and the ankle (right)

Table 2 Quality metric values

Dataset	Sample fraction	Metric	PICS	PICS+SR	PICS+Ω	PICS+SR+Ω
Brain	0.35	SSIM	0.966	0.967	0.967	0.967
	0.30		0.962	0.964	0.963	0.964
	0.25		0.959	0.959	0.959	0.959
	0.20		0.952	0.952	0.952	0.952
	0.35	PSNR (dB)	40.15	40.24	40.15	40.24
	0.30		39.60	39.69	39.60	39.69
	0.25		39.06	39.15	39.06	39.15
	0.20		38.38	38.47	38.38	38.47
Knee	0.35	SSIM	0.800	0.814	0.800	0.814
	0.30		0.779	0.796	0.779	0.7964
	0.25		0.750	0.772	0.750	0.772
	0.20		0.721	0.747	0.721	0.746
	0.35	PSNR (dB)	29.66	30.42	29.66	30.42
	0.30		29.12	29.93	29.12	29.92
	0.25		28.49	29.30	28.49	29.30
	0.20		27.89	28.67	27.89	28.66
Ankle	0.35	SSIM	0.974	0.975	0.974	0.975
	0.30		0.970	0.971	0.970	0.971
	0.25		0.966	0.967	0.966	0.967
	0.20		0.960	0.961	0.960	0.961
	0.35	PSNR (dB)	40.74	40.83	40.74	40.83
	0.30		39.94	40.03	39.94	40.03
	0.25		39.06	39.16	39.06	39.16
	0.20		37.95	38.06	37.95	38.06

Quality metric values for reconstructions using parallel imaging and compressed sensing (PICS), PICS with SPIRiT regularization (SR), PICS with the support (Ω), and PICS with SPIRiT regularization as well as the support. The best reconstruction with respect to each metric is underlined

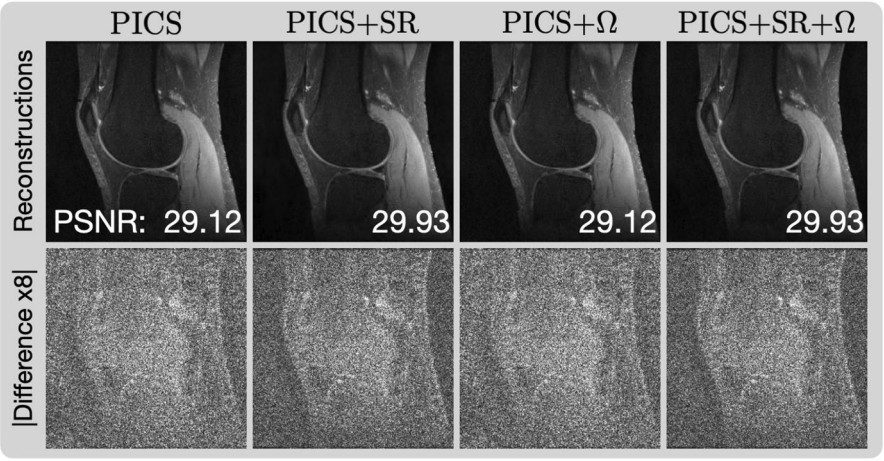

Fig. 7 Reconstructions of the knee with a 0.3 sample fraction using parallel imaging and compressed sensing (PICS), PICS with SPIRiT regularization (SR), PICS with the support (Ω), and PICS with SPIRiT regularization as well as the support. The second row shows the magnitudes of the differences between the under-sampled reconstructions and the fully-sampled reconstructions multiplied by 8. The PSNR for each reconstruction is also presented

However, this would likely include an additional parameter for the threshold at which to stop the iterations. Again, we will investigate this approach in future work.

Other future efforts will alter the optimization problem presented in this manuscript to incorporate partial Fourier sampling [3], similar to [4], and structured sparsity [41, 44, 45].

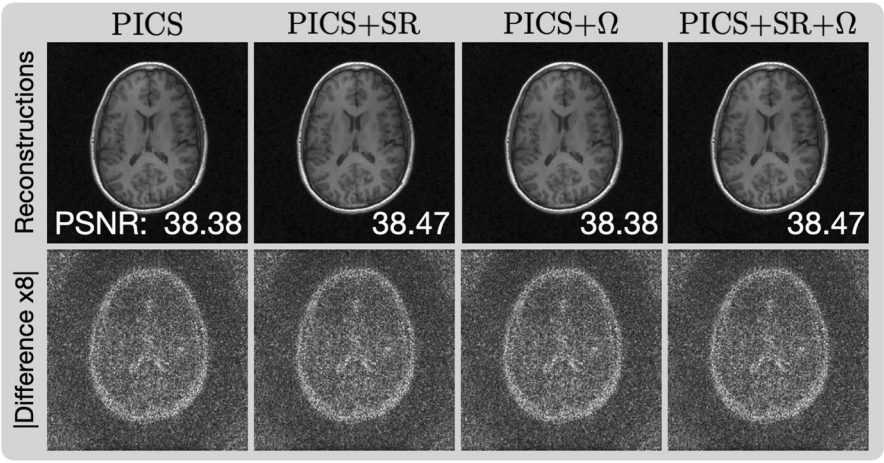

Fig. 8 Reconstructions of the brain with a 0.2 sample fraction using parallel imaging and compressed sensing (PICS), PICS with SPIRiT regularization (SR), PICS with the support (Ω), and PICS with SPIRiT regularization as well as the support. The second row shows the magnitudes of the differences between the under-sampled reconstructions and the fully-sampled reconstructions multiplied by 8. The PSNR for each reconstruction is also presented in decibels

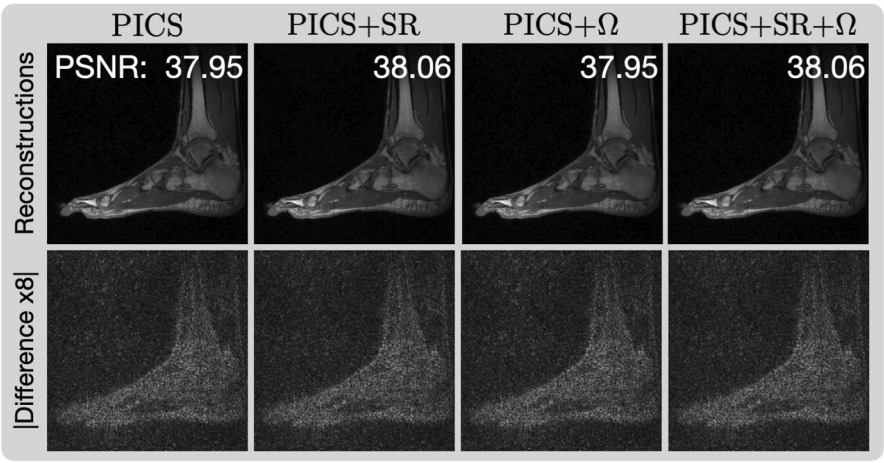

Fig. 9 Reconstructions of the ankle with a 0.2 sample fraction using parallel imaging and compressed sensing (PICS), PICS with SPIRiT regularization (SR), PICS with the support (Ω), and PICS with SPIRiT regularization as well as the support. The second row shows the magnitudes of the differences between the under-sampled reconstructions and the fully-sampled reconstructions multiplied by 8. The PSNR for each reconstruction is also presented in decibels

Discrete Fourier Transform

The definition of the two-dimensional discrete Fourier transform F used in this manuscript is

$$F\,x = \sum_{u=0}^{U-1} \sum_{v=0}^{V-1} x_{u,v} \exp\left(-i\,2\pi\left(\frac{k_u\,u}{U} + \frac{k_v\,v}{V}\right)\right).$$

The inverse Discrete Fourier transform is

$$F^{-1}\,y = \frac{1}{UV} \sum_{k_u=0}^{U-1} \sum_{k_v=0}^{V-1} \hat{x}_{u,v} \exp\left(i\,2\pi\left(\frac{k_u\,u}{U} + \frac{k_v\,v}{V}\right)\right),$$

where $\hat{x} = F\,x$.

Inclusion of the Support

The process of clinical MRI consists of 1) reconstructing a low-resolution *localizer* image to ensure the patient is properly positioned in the bore and from which the technologist selects the field-of-view ([FOV] - the volume of space to image with high fidelity), and 2) reconstruct a high-quality image for the volume within the FOV from a second set of collected data. Conventionally, the FOV selected by the technologist has been a rectangle. However, the system could be altered so that the technologist could indicate the support of the object to be imaged (i.e., a

non-rectangular region that contains the object). The FOV would then be the smallest rectangle that contains this support. Alternatively, an automatic system—such as PISCO [24]—could identify the non-rectangular support from the data used to reconstruct the localizer image. The support, which encloses a smaller volume than the FOV, could require less data to reconstruct a high-quality image [14].

Let the support of the image be the set of those locations in space where the true image (without noise) is non-zero. Let Ω be a superset of the support. In past works, this superset has been used to improve image quality for a given number of samples [12]. Let $T_{\bar{\Omega}}$ be the linear transformation that extracts those pixels outside of the support Ω and arranges them into a vector. (Here, $\bar{\Omega}$ is the complement of the set of pixels within the support.) Then, problem (6) is modified to include the support as follows.

$$\begin{aligned} \underset{m}{\text{minimize}} \;\; & (1/2)\|\boldsymbol{M}\boldsymbol{F}\boldsymbol{S}\,m-\boldsymbol{b}\|_2^2 \quad + \nu\|\Psi m\|_1 + (\lambda_s/2)\sum_{c=1}^{C}\| \\ & \left(W^{(c)}-I\right)F\,S^{(c)}\,m\Big\|_{\gamma,2}^2 \quad \text{subject to} \quad \frac{\|T_{\bar{\Omega}}\,m\|_2^2}{|\bar{\Omega}|} \le \sigma^2. \end{aligned} \tag{9}$$

Here, $|\bar{\Omega}|$ is the number of elements in $\bar{\Omega}$, and σ^2 is a bound on the average energy for those pixels outside of the support. Unlike the regularization parameters ν and λ_s, σ^2 can be estimated from data by computing the sample variance of an image reconstructed from data captured during an acquisition without any excitation.

The optimization problem of (9) can be solved with the primal-dual hybrid gradient (PDHG) method [48]. Specifically, we will use PDHG with line search, as detailed in [49]. PDHG solves problems of the form

$$\text{minimize} \;\; V(y) + W(A\,y), \tag{10}$$

where V and W are both proper, convex, and lower semicontinuous (PCLSC) with computationally feasible proximal operators and A is a linear transformation.

To put the problem (9) is put in the form of Eq. (10), first note that Eq. (9) is equivalent to

$$\begin{aligned} \underset{m}{\text{minimize}} \;\; & \frac{1}{2}\,\|\boldsymbol{M}\boldsymbol{F}\boldsymbol{S}\,m-\boldsymbol{b}\|_2^2 + (\lambda_s/2)\sum_{c=1}^{C}\| \\ & \left(W^{(c)}-I\right)F\,S^{(c)}\,m\Big\|_{\gamma,2}^2 + \nu\|\Psi m\|_1 + \mathbb{I}_{B\left[0,\sigma\sqrt{|\Omega|}\right]}(P_{\bar{\Omega}}\,m), \end{aligned} \tag{11}$$

where $B\left[0,\sqrt{\nu\,|\Omega|}\right]$ is the closed ball centered on 0 with radius $\sqrt{\nu\,|\Omega|}$. The function $\mathbb{I}_Q(\cdot)$ is the indicator function for the set Q: equal to 0 when its argument is within the set and otherwise equal to infinity.

Let $V(m) = \|\Psi m\|_1$,

$$A = \begin{bmatrix} \boldsymbol{MFS} \\ \left(W^{(1)} - I\right) F S^{(1)} \\ \left(W^{(2)} - I\right) F S^{(2)} \\ \vdots \\ \left(W^{(C)} - I\right) F S^{(C)} \\ I \end{bmatrix}$$

and

$$W(y^{(1)}, y^{(2)}, \ldots, y^{(C+1)}, y^{(C+2)}) = \frac{1}{2}\|y^{(1)} - \boldsymbol{b}\|_2^2 + \frac{\lambda_s}{2} \sum_{c=1}^{C} \left\| y^{(c+1)} \right\|_{\gamma,2}^2 + \mathbb{I}_{B\left[0, \sqrt{v\,|\Omega|}\right]}(P_{\bar{\Omega}}\, m).$$

It has already been established that V is PCLSC with a simple proximal operator. Similarly, W is PCLSC. Since each term in the sum of W has a simple proximal operator, and since each term operates on different components of y, W has a simple proximal operator by the separable sum property of proximal operators. Therefore, problem (10) is equivalent to Eq. (11) which can be solved with PDHG (See Table 2).

For this work, since data collected without excitation was unavailable, σ^2 was estimated using the sample variance of a region of the fully-sampled reconstruction that was outside of the anatomy. For this work, the support regions were identified with PISCO [24]; they are shown in Fig. 6.

In all cases, the difference in metric value when including the support is extremely small (and perhaps negligible). With respect to PSNR, the inclusion of the support improves the outcome in the majority of reconstructions. With respect to SSIM, though, the inclusion of the support yields the best metric value in only half of the reconstructions.

Figure 7 shows reconstructions of the knee with a sample fraction of 0.3, and figures 8, and 9 show the reconstructions of the brain and ankle with a 0.20 sample fraction for the four different reconstruction methods along with the magnitudes of the differences from the fully-sampled images multiplied by 8. The difference images indicate that the improvement in image quality when including the support is negligible.

Note that problem (9) contains a differentiable term, making it amenable to being solved with the Condat-Vu algorithm [50] or with PD3O [51]. Should a sampling pattern or support region selection system be discovered that improves image quality when incorporating the support, these alternatives to PDHG may converge to a high-quality solution faster.

Author Contributions All authors reviewed the manuscript. N.D. conceived of the approach, implemented the methods, and designed the experiments. This work was based on that of A.M., who implemented the earlier method and provided his thoughts during development. S.B. was an expert of optimization and an advisor who guided the analysis, reviewed the material, and provided advice on how to proceed with the optimization. G.S. conducted all experiments and analysis.

Funding This work was supported in part by a Rifkin and Bennis Cancer Imaging AI Early-stage Project Award. This work was supported by the University of Colorado through a research grant awarded by the Anschutz-Boulder Nexus (AB Nexus) Initiative.

Data Availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no conflict of interest.

Ethical Approval This study did not involve human subjects or animals, and, therefore, no ethical approval or informed consent was required. The research complies with all relevant ethical guidelines.

References

1. D.C. Noll, D.G. Nishimura, A. Macovski, Homodyne detection in magnetic resonance imaging. IEEE Trans. Med. Imaging **10**, 154–163 (1991)
2. G. McGibney, M. Smith, S. Nichols, A. Crawley, Quantitative evaluation of several partial fourier reconstruction algorithms used in MRI. Magn. Reson. Med. **30**, 51–59 (1993)
3. M. Bydder, M.D. Robson, Partial Fourier partially parallel imaging. Magn. Reson. Med. **53**, 1393–1401 (2005)
4. A. McManus, S.R. Becker, D. O'Connor, N. Dwork, Compressed sensing with homodyne detection. IEEE Conference on Computational Imaging Using Synthetic Apertures (CISA) 01–05 (2024)
5. A. Deshmane, V. Gulani, M.A. Griswold, N. Seiberlich, Parallel MR imaging. J. Magn. Reson. Imaging **36**, 55–72 (2012)
6. K.P. Pruessmann, M. Weiger, P. Börnert, P. Boesiger, Advances in sensitivity encoding with arbitrary k-space trajectories. Magn. Reson. Med. **46**, 638–651 (2001)
7. K.P. Pruessmann, Encoding and reconstruction in parallel MRI. NMR Biomed. **19**, 288–299 (2006)
8. D.L. Donoho, M. Elad, V.N. Temlyakov, Stable recovery of sparse overcomplete representations in the presence of noise. IEEE Trans. Inf. Theory **52**, 6–18 (2005)
9. M. Lustig, D. Donoho, J.M. Pauly, Sparse MRI: The application of compressed sensing for rapid MR imaging. Magn. Reson. Med. **58**, 1182–1195 (2007)
10. E.J. Candès, M.B. Wakin, An introduction to compressive sampling. IEEE Signal Process. Mag. **25**, 21–30 (2008)
11. S.K. Plevritis, A. Macovski, Spectral extrapolation of spatially bounded images. IEEE Trans. Med. Imaging **14**, 487–497 (1995)
12. A.A. Samsonov, E.G. Kholmovski, D.L. Parker, C.R. Johnson, POCSENSE: POCS-based reconstruction for sensitivity encoded magnetic resonance imaging. Magn. Reson. Med. **52**, 1397–1406 (2004)
13. X. Peng, H.-Q. Du, F. Lam, S.D. Babacan, Z.-P. Liang, Reference-driven MR image reconstruction with sparsity and support constraints. IEEE International Symposium on Biomedical Imaging 89–92 (2011)
14. N. Dwork, E.K. Englund, A.J. Barker, Reducing the sampling burden of Fourier sensing with a non-rectangular field-of-view. Appl. Math. Mod. Challenges **4**, 1–16 (2025)
15. D.K. Sodickson, W.J. Manning, Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. Magn. Reson. Med. **38**, 591–603 (1997)
16. M.A. Griswold et al., Generalized autocalibrating partially parallel acquisitions (GRAPPA). Magn. Reson. Med. **47**, 1202–1210 (2002)
17. J.P. Haldar, K. Setsompop, Linear predictability in magnetic resonance imaging reconstruction: Leveraging shift-invariant fourier structure for faster and better imaging. IEEE Signal Process. Mag. **37**, 69–82 (2020)
18. A. McManus, S. Becker, N. Dwork, Dependence of parallel imaging with linear predictability on the undersampling direction. J. Electron. Imaging **34**, 023031–023031 (2025)
19. K.P. Pruessmann, M. Weiger, M.B. Scheidegger, P. Boesiger, SENSE: sensitivity encoding for fast MRI. Magn. Reson. Med. **42**, 952–962 (1999)

20. K.P. Pruessmann, M. Weiger, P. Boesiger, Sensitivity encoded cardiac MRI. J. Cardiovasc. Magn. Reson. **3**, 1–9 (2001)
21. H.C.M. Holme et al., ENLIVE: an efficient nonlinear method for calibrationless and robust parallel imaging. Sci. Rep. **9**, 3034 (2019)
22. J.A. Fessler, Optimization methods for magnetic resonance image reconstruction: Key models and optimization algorithms. IEEE Signal Process. Mag. **37**, 33–40 (2020)
23. M. Uecker et al., ESPIRiT—an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA. Magn. Reson. Med. **71**, 990–1001 (2014)
24. R.A. Lobos, C.-C. Chan, J.P. Haldar, New theory and faster computations for subspace-based sensitivity map estimation in multichannel MRI. IEEE Trans. Med. Imaging **43**, 286–296 (2023)
25. M. Uecker, J.I. Tamir, F. Ong, M. Lustig, The BART toolbox for computational magnetic resonance imaging. Proc. Int. Soc. Magn. Reson. Med. **24**, 1 (2016)
26. P. Beatty et al., A method for autocalibrating 2D accelerated volumetric parallel imaging with clinically practical reconstruction times. Proc. Int. Soc. Magn. Reson. Med. **15**, 1749 (2007)
27. M. Lustig, J.M. Pauly, SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space. Magn. Reson. Med. **64**, 457–471 (2010)
28. D. Liang, B. Liu, J. Wang, L. Ying, Accelerating SENSE using compressed sensing. Magn. Reson. Med. Off. J. Int. Soc. Magn. Reson. Med. **62**, 1574–1584 (2009)
29. A. Majumdar, R.K. Ward, On the choice of compressed sensing priors and sparsifying transforms for mr image reconstruction: an experimental study. Signal Process. Image Commun. **27**, 1035–1048 (2012)
30. A. Dutt, V. Rokhlin, Fast Fourier transforms for nonequispaced data. SIAM J. Sci. Comput. **14**, 1368–1393 (1993)
31. A. Dutt, V. Rokhlin, Fast Fourier transforms for nonequispaced data. II. Appl. Comput. Harmon. Anal. **2**, 85–100 (1995)
32. N. Dwork et al., Optimization in the space domain for density compensation with the nonuniform FFT. Magn. Reson. Imaging **100**, 102–111 (2023)
33. R.N. Bracewell, *Two-dimensional imaging* (Prentice-Hall Inc, New Jersey, 1995)
34. D.L. Ruderman, Origins of scaling in natural images. Vision. Res. **37**, 3385–3398 (1997)
35. statistics and information, Van der Schaaf, v. A. & van Hateren, J. v. Modelling the power spectra of natural images. Vision. Res. **36**, 2759–2770 (1996)
36. S. Boyd, L. Vandenberghe, *Introduction to applied linear algebra: vectors, matrices, and least squares* (Cambridge University Press, Cambridge, 2018)
37. A. Beck, M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM J. Imag. Sci. **2**, 183–202 (2009)
38. D. O'Connor, V. Yu, D. Nguyen, D. Ruan, K. Sheng, Fraction-variant beam orientation optimization for non-coplanar IMRT. Phys. Med. Biol. **63**, 045015 (2018)
39. P.L. Combettes, V.R. Wajs, Signal recovery by proximal forward-backward splitting. Multisc. Model. Simulat. **4**, 1168–1200 (2005)
40. Ong, F., Amin, S., Vasanawala, S. & Lustig, M. Mridata.org: An open archive for sharing MRI raw data. *Proceedings of the International Society for Magnetic Resonance in Medicine***26** (2018). www.mridata.org
41. N. Dwork, J.W. Gordon, E.K. Englund, Accelerated parallel magnetic resonance imaging with compressed sensing using structured sparsity. J. Med. Imaging **11**, 033504–033504 (2024)
42. N. Dwork et al., Fast variable density Poisson-disc sample generation with directional variation for compressed sensing in MRI. Magn. Reson. Imaging **77**, 186–193 (2021)
43. B. Adcock, A.C. Hansen, C. Poon, B. Roman, Breaking the coherence barrier: a new theory for compressed sensing. Forum Math. Sigma **5**, e4 (2017)
44. N. Dwork et al., Utilizing the wavelet transform's structure in compressed sensing. SIViP **15**, 1407–1414 (2021)
45. N. Dwork, P.E. Larson, Utilizing the structure of a redundant dictionary comprised of wavelets and curvelets with compressed sensing. J. Electron. Imaging **31**, 063043–063043 (2022)
46. Z. Wang, A.C. Bovik, H.R. Sheikh, E.P. Simoncelli, Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. **13**, 600–612 (2004)
47. F. Ong, Low dimensional methods for high dimensional magnetic resonance imaging (Berkeley, 2018), (Ph.D. thesis)
48. A. Chambolle, T. Pock, A first-order primal-dual algorithm for convex problems with applications to imaging. J. Math. Imaging Vis. **40**, 120–145 (2011)

49. Y. Malitsky, T. Pock, A first-order primal-dual algorithm with linesearch. SIAM J. Optim. **28**, 411–432 (2018)
50. L. Condat, A primal-dual splitting method for convex optimization involving lipschitzian, proximable and linear composite terms. J. Optim. Theory Appl. **158**, 460–479 (2013)
51. M. Yan, *A primal-dual three-operator splitting scheme* (Tech. Rep, Technical Report, 2016)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.