

Projection onto Epigraph Sets for Rapid Self-Tuning Compressed Sensing MRI

Mohammad Shahdloo, Efe Ilicak, Mohammad Tofghi, Emine U. Saritas, *Member, IEEE*,
A. Enis Çetin, *Fellow, IEEE*, and Tolga Çukur*, *Senior Member, IEEE*

Abstract—The compressed sensing (CS) framework leverages the sparsity of MR images to reconstruct from undersampled acquisitions. CS reconstructions involve one or more regularization parameters that weigh sparsity in transform domains against fidelity to acquired data. While parameter selection is critical for reconstruction quality, the optimal parameters are subject and dataset specific. Thus, commonly practiced heuristic parameter selection generalizes poorly to independent datasets. Recent studies have proposed to tune parameters by estimating the risk of removing significant image coefficients. Line searches are performed across the parameter space to identify the parameter value that minimizes this risk. Although effective, these line searches yield prolonged reconstruction times. Here, we propose a new self-tuning CS method that uses computationally efficient projections onto epigraph sets of the ℓ_1 and total-variation norms to simultaneously achieve parameter selection and regularization. In vivo demonstrations are provided for balanced steady-state free precession, time-of-flight, and T1-weighted imaging. The proposed method achieves an order of magnitude improvement in computational efficiency over line-search methods while maintaining near-optimal parameter selection.

Index Terms—Compressed sensing (CS), magnetic resonance imaging (MRI), projection onto epigraph sets, self-tuning, parameter selection, multi-coil, multi-acquisition.

I. INTRODUCTION

THE compressed sensing (CS) framework was recently proposed for accelerated MRI, where compressibility of MR images are employed to reconstruct from undersampled acquisitions [1]–[4]. To do this, CS reconstructions are typically cast as regularized optimization problems that weigh data consistency against sparsity in some transform domain (e.g., wavelet domain, total variation (TV)) [1]. The weighing

between data consistency and sparsity is governed by regularization parameters. High parameter values overemphasize sparsity at the expense of introducing inconsistency to acquired data samples, potentially leading to feature losses. Meanwhile, low parameter values render the reconstructions ineffective in suppressing residual aliasing and noise in undersampled acquisitions. Since the optimal regularization parameters are subject and dataset specific, time-consuming and potentially erroneous heuristic selection is performed in many studies, limiting the clinical utility of CS-MRI.

Several unsupervised methods have been proposed to address parameter selection in CS-MRI. Empirical methods including the L-curve criterion (LCC) follow the notion that the optimal parameter should be selected to attain a favorable trade-off between data consistency and regularization objectives [5]–[7]. Assuming this trade-off is characterized by an L-shaped curve, LCC selects the parameter on the point of maximum curvature [8]. LCC has been successfully demonstrated for parameter selection in several applications including parallel imaging [9], [10], quantitative susceptibility mapping [11], and diffusion spectrum imaging [12]. However, curvature assessment is computationally inefficient and typically sensitive to numerical perturbation and nonlinearities in the reconstruction problem [13]–[15].

Alternatively, parameters can be selected based on analytical estimates of the reconstruction error to optimize the regularization parameters. These methods include generalized cross-validation (GCV) [16], and methods based on Stein’s unbiased risk estimator (SURE) [8], [17]. In GCV, an analytical measure for reconstruction error is estimated that asymptotically converges to the true error [16]. The GCV measure is derived as a function of the sampling pattern, regularization function, and regularization parameter. Parameter estimation via minimization of the GCV measure has been used in a variety of applications such as functional MRI [18], perfusion imaging [19], and dynamic MRI [20]. However, the GCV measure can be expensive to compute and yields biased estimates of the true error with limited number of data samples [21].

A recent approach instead uses the SURE criterion to estimate the expected value of the mean-square error (MSE) of the reconstruction. Given a specific parameter value and an estimate of the noise variance, Stein’s lemma [17] is used to compute online estimates of MSE. Subsequently, a line search over potential parameter values is performed for selecting the optimal parameter at each iteration. SURE-based parameter selection has produced promising results in several sparse recovery applications including CS-MRI [22]–[27].

This work was supported in part by TUBITAK 1001 Grant (117E171), by a European Molecular Biology Organization Installation Grant (IG 3028), by a TUBA GEBIP fellowship, and by BAGEP 2016 and BAGEP 2017 awards of the Science Academy. *Asterisk indicates the corresponding author.*

M. Shahdloo, E. Ilicak, E. U. Saritas, and T. Çukur are with the Department of Electrical and Electronics Engineering and the National Magnetic Resonance Research Center, Bilkent University, Bilkent, 06800 Ankara, Turkey (e-mail: shahdloo@ee.bilkent.edu.tr; ilicak@ee.bilkent.edu.tr; saritas@ee.bilkent.edu.tr; cukur@ee.bilkent.edu.tr).

A. E. Çetin is with the Department of Electrical and Computer Engineering, University of Illinois at Chicago, 1020 Sciences and Engineering Offices (SEO), 851 South Morgan St.(M/C 154) Chicago, IL 60607, on leave from the Department of Electrical and Electronics Engineering, Bilkent University, Ankara, Turkey (e-mail: cetin@bilkent.edu.tr)

M. Tofghi is with the Department of Electrical Engineering, Pennsylvania State University, 104 EE East Building, State College, PA (email: tofghi@psu.edu)

Copyright © 2018 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

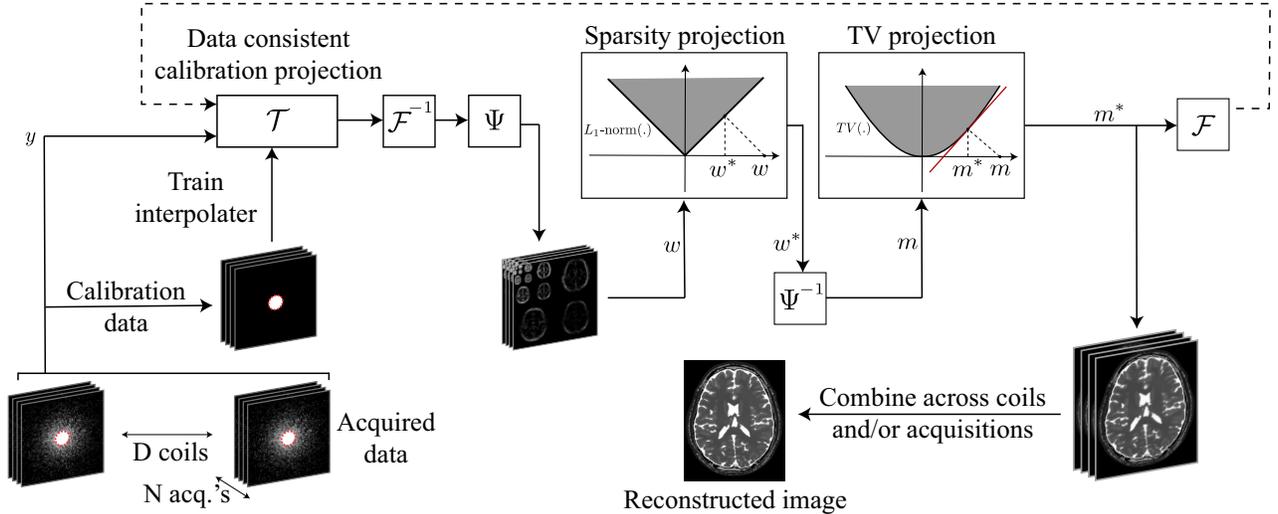


Fig. 1. Flowchart of the PESCaT reconstruction. PESCaT employs an alternating projections onto sets approach with three subprojections: data-consistent calibration projection, sparsity projection, and TV projection. The calibration projection linearly synthesizes unacquired k-space samples via a tensor interpolating kernel. The sparsity projection jointly projects wavelet coefficients of the multi-coil, multi-acquisition images onto the epigraph set of the ℓ_1 -norm function. The TV projection projects image coefficients onto the epigraph set of the TV-norm function. These projections are performed iteratively until convergence. Lastly, reconstructed images are combined across multiple coils and/or acquisitions.

Unfortunately, parameter searches that need to be performed in each iteration cause substantial computational burden.

Here we introduce a computationally efficient self-tuning reconstruction method, named PESCaT (Projection onto Epigraph Sets for reconstruction by Calibration over Tensors), that can handle both single-acquisition and multi-acquisition datasets. To jointly reconstruct undersampled acquisitions, PESCaT performs tensor-based interpolation across acquired data, complemented by sparsity regularization of wavelet coefficients and TV regularization of image coefficients.

Since wavelet coefficients show varying sparsity across subbands and decomposition levels, PESCaT uses different ℓ_1 regularization parameters for each subband and level. Similarly, multi-coil multi-acquisition image coefficients may show varying spatial gradients, so different TV regularization parameters are used for each coil and acquisition. Parameters are efficiently tuned via simple geometric projections onto the boundary of the convex epigraph sets for the ℓ_1 - and TV-norm functions. This formulation transforms the selection of many different regularization parameters for multiple subbands, levels, coils, and acquisitions into the selection of two scaling factors for the ℓ_1 -norm and TV-norm epigraphs. These factors can be reliably tuned on training data, yielding consistent performance across sequences, acceleration factors, and subjects. Comprehensive demonstrations on simulated brain phantoms, and in vivo balanced steady-state free-precession (bSSFP), T1-weighted, and angiographic acquisitions indicate that PESCaT enables nearly an order of magnitude improvement in computational efficiency compared to SURE-based methods, without compromising reconstruction quality.

II. THEORY

Our main aim is to introduce a fast joint reconstruction method that automatically selects the free parameters for regularization terms based on ℓ_1 - and TV-norms. We consider the application of this self-tuning reconstruction to single-coil multi-acquisition, multi-coil single-acquisition, and multi-

coil multi-acquisition MRI datasets. In the following sections, we introduce the regularized reconstruction problem, and its solution via projection onto epigraph sets for unsupervised parameter selection.

A. Reconstruction by calibration over tensors

Compressive sensing (CS) techniques proposed for static MRI acquisitions typically leverage encoding information provided either by multiple coils [2], [28], [29] or by multiple acquisitions [30]–[32] to enable recovery of unacquired data samples. Yet, simultaneous use of information across coils and acquisitions can benefit phase-cycled bSSFP [33], [34], multi-contrast [35]–[37] or parametric imaging [38]–[40]. Here we consider a joint reconstruction framework for multi-coil, multi-acquisition datasets, based on a recent method that we have proposed named ReCaT (Reconstruction by Calibration over Tensors) [33]. ReCaT rests on the following spatial encoding model for the signal measured in acquisition $n \in [1, \dots, N]$ and coil $d \in [1, \dots, D]$:

$$S_{nd}(r) = P_n(r)C_d(r)S_0(r), \quad (1)$$

where r is the spatial location, P_n is the acquisition spatial profile, C_d is the coil sensitivity profile, and S_0 is the signal devoid of coil sensitivity and acquisition profile modulations. ReCaT seeks to linearly synthesize missing k-space samples from neighboring acquired samples across all coils and acquisitions. A tensor interpolation kernel is used for this purpose:

$$x_{nd} = \sum_{i=1}^N \sum_{j=1}^D t_{ij,nd}(k_r) \otimes x_{ij}(k_r), \quad (2)$$

where x_{nd} is the k-space data from n^{th} acquisition and d^{th} coil, k_r is the k-space location, and \otimes is the convolution operation. Here $t_{ij,nd}(k_r)$ accounts for the contribution of samples from acquisition i and coil j to x_{nd} . Equation (2) can be compactly expressed as:

$$x = \mathcal{T}x. \quad (3)$$

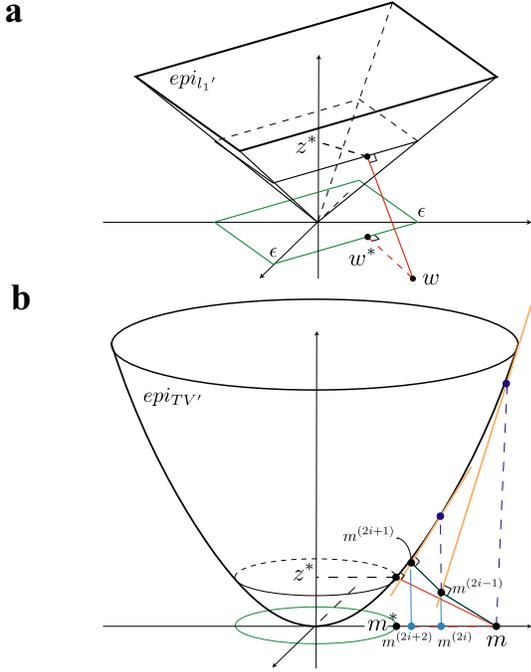


Fig. 2. The projection onto epigraph sets (PES) approach illustrated in \mathbb{R}^3 . **(a)** PES for ℓ_1 -regularization. An input vector w (e.g., vector of wavelet coefficients of image m) is projected onto the epigraph set of the ℓ_1 -norm function (epi_{ℓ_1}). This projection results in the output $[w^*z^*]^T$, thereby inherently calculating the projection of w onto the ℓ_1 -ball in \mathbb{R}^2 (w^*). The size of the ℓ_1 -ball (ϵ) depicted in green color depends on z^* . **(b)** PES for TV regularization. Unlike PES- ℓ_1 , PES-TV has no closed-form solution, and is instead implemented via an iterative epigraphical splitting procedure. At the i^{th} iteration, the input vector $m^{(2i-1)}$ is projected onto the supporting hyperplane (orange line), resulting in $m^{(2i+1)}$. This intermediate vector is then projected on the level set to compute $m^{(2i+2)}$. Through successive iterations the output gradually converges to the desired projection point on the epigraph set $[m^*z^*]^T$, thereby inherently calculating the projection of m onto the TV-ball in \mathbb{R}^2 (m^*). The size of the TV-ball depends on z^* .

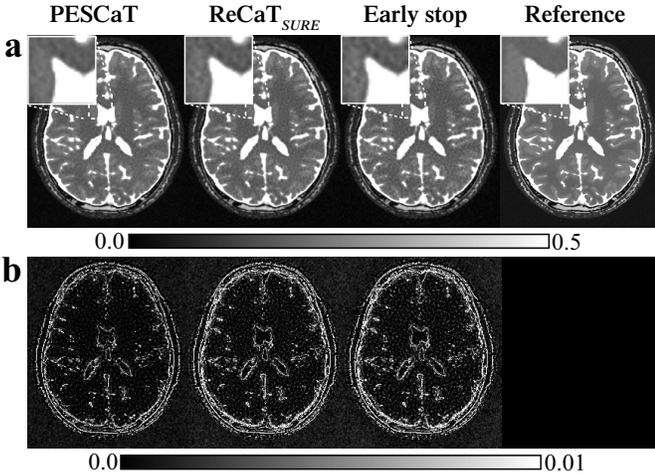


Fig. 3. Reconstructions of phase-cycled bSSFP acquisitions of the simulated brain phantom. Simulations assumed $\alpha=45^\circ$, $\text{TR}/\text{TE}=5/2.5$ ms, a fieldmap of 0 ± 62 Hz (mean \pm std), $N = 6$ phase-cycles and $R = 6$. **(a)** PESCaT, ReCaT_{SURE}, ReCaT_{SURE} with early stop, and the fully-sampled reference images are shown. White boxes display zoomed-in portions of the images. **(b)** Mean-squared error between the reconstructed and reference images are shown. PESCaT yields visibly reduced errors compared to both ReCaT_{SURE} and ReCaT_{SURE} with early stop.

B. PESCaT

In the previous study where we proposed ReCaT, a basic implementation was considered that did not include any reg-

ularization terms to enforce sparsity [33]. Here we introduce an improved version, PESCaT, that incorporates sparsity and TV penalties:

$$\begin{aligned} \min_{x_{nd}} \left\{ \sum_{n=1}^N \sum_{d=1}^D \|(\mathcal{T} - \mathcal{I})x_{nd}\|_2^2 \right. \\ + \sum_{l=1}^L \sum_{s=1}^3 \lambda_{\ell_1,ls} \sum_{n=1}^N \sum_{d=1}^D \|\Psi_{ls}\{\mathcal{F}^{-1}\{x_{nd}\}\}\|_1 \\ \left. + \sum_{n=1}^N \sum_{d=1}^D \lambda_{TV,nd} \|\mathcal{F}^{-1}\{x_{nd}\}\|_{TV} \right\}, \end{aligned} \quad (4)$$

where Ψ_{ls} is the wavelet operator for subband s and level l , \mathcal{I} is the identity operator, and \mathcal{F}^{-1} is the inverse Fourier operator. A separate ℓ_1 -regularization parameter, $\lambda_{\ell_1,ls}$, is prescribed for each subband and level of the wavelet coefficients. Sparsity regularization is performed on the three high-pass subbands while the low-pass subband is kept intact to avoid over-smoothing. Meanwhile, a separate TV regularization parameter, $\lambda_{TV,nd}$, is used for each acquisition and coil. Because wavelet coefficients are aggregated across the coil and acquisition dimensions, $\lambda_{\ell_1,ls}$ varies across wavelet levels and subbands but it is fixed across coils or acquisitions.

In this study, we implemented PESCaT in a constrained optimization formulation equivalent to the Lagrangian formulation in (4):

$$\begin{aligned} \min_{x_{nd}} \quad & \sum_{n=1}^N \sum_{d=1}^D \|(\mathcal{T} - \mathcal{I})x_{nd}\|_2^2 \\ \text{subject to} \quad & \sum_{n=1}^N \sum_{d=1}^D \|\Psi_{ls}\{\mathcal{F}^{-1}\{x_{nd}\}\}\|_1 \leq \epsilon_{\ell_1,ls} \\ & s = 1, 2, 3 \\ & l = 1, \dots, L; \\ & \|\mathcal{F}^{-1}\{x_{nd}\}\|_{TV} \leq \epsilon_{TV,nd} \\ & n = 1, \dots, N \\ & d = 1, \dots, D \end{aligned} \quad (5)$$

where $\epsilon_{\ell_1,ls}$ are the constraints on the sparsity of the reconstruction, and $\epsilon_{TV,nd}$ are the constraints on the TV of the reconstruction. The optimization problem in (5) was solved via an alternating projections onto sets algorithm. As outlined in Fig. 1, this algorithm involves three consecutive projections, namely data-consistent calibration, sparsity, and TV projections. The calibration projection linearly synthesized unacquired k-space samples via the tensor interpolating kernel. To perform this projection while enforcing strict consistency to acquired data, an iterative least-squares algorithm was employed [28]. The sparsity projection jointly projected wavelet coefficients of images onto the epigraph set of the ℓ_1 -norm function. The TV projection projected image coefficients onto the epigraph set of the TV-norm function. These projections were performed iteratively until convergence. At each iteration, MSE between the reconstructed image in the current iterate and the previous iterate was first measured, and the percentage change in MSE across consecutive iterations was then calculated. Convergence was taken to be the iteration

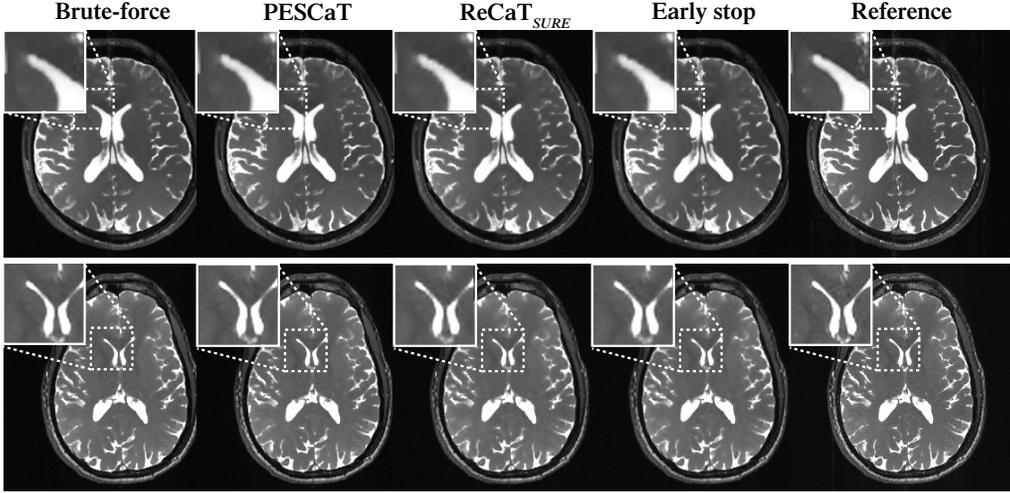


Fig. 4. Reconstructions of in vivo bSSFP acquisitions of the brain at $R=6$. Brute-force, PESCaT, $\text{ReCaT}_{\text{SURE}}$, $\text{ReCaT}_{\text{SURE}}$ with early stop, and reference images are shown in two representative subjects. White boxes display zoomed-in portions of the images. PESCaT achieves significantly improved image quality compared to $\text{ReCaT}_{\text{SURE}}$ with early stop that was matched PESCaT in terms of the total reconstruction time. Furthermore, PESCaT yields similar image quality to $\text{ReCaT}_{\text{SURE}}$ and brute-force methods, while also maintaining greater computational efficiency.

at which the percentage change in MSE fell below 20%. Lastly, reconstructed images were combined across multiple coils and/or acquisitions. Note that because PESCaT is structured modularly regarding the calibration, sparsity, and TV projections, it is trivial to implement variants that only employ sparsity or TV regularization.

C. Parameter tuning by projection onto epigraph sets

Careful tuning of constraint parameters in (5) is critical for a successful reconstruction. Selecting too tight constraints can lead to loss of important image features, whereas selecting too loose constraints will yield substantial residual noise and aliasing. When only a few parameters are to be tuned, an exhaustive search over a relevant range of values followed by visual inspection is typically exercised. However, even in a modest dataset with $D = 4$ coils and $N = 4$ acquisitions, and assuming $L = 4$ wavelet decomposition levels there are 28 distinct parameters involved in (5). Thus, the exhaustive search approach is impractical.

In this study, we perform self-tuning of the constraint parameters in (5) via projections onto epigraph sets of the respective regularization terms. Let $\mathcal{U} \in \mathbb{R}^k$ be a closed convex set, $\Phi : \mathbb{R}^k \rightarrow \mathbb{R}$ be a convex function (e.g., ℓ_1 -norm and TV-norm functions), and $\hat{u} \in \mathbb{R}^k$ be an input vector (e.g., wavelet coefficients for ℓ_1 -norm or image coefficients for TV-norm). The proximal operator of Φ^2 is:

$$\text{prox}_{\Phi^2}(\hat{u}) = \arg \min_{u \in \mathcal{U}} \|\hat{u} - u\|_2^2 + \Phi^2(u), \quad (6)$$

where u is the auxiliary variable. We prefer to use Φ^2 here since it allows us to express the solution as a simple geometric projection. Specifically, the problem in (6) can be stated in vector form by mapping onto \mathbb{R}^{k+1} :

$$\min_{u \in \mathcal{U}} \left\| \begin{bmatrix} \hat{u} \\ 0 \end{bmatrix} - \begin{bmatrix} u \\ \Phi(u) \end{bmatrix} \right\|_2^2. \quad (7)$$

Here we propose to implement the proximal operator in (6) by identifying the closest vector $[u^* \ \Phi(u^*)]^T \in \mathbb{R}^{k+1}$ to $[\hat{u} \ 0]^T$. This solution can be shown to be equivalent to the

orthogonal projection of the vector $[\hat{u} \ 0]^T$ onto the epigraph set of Φ (epi_{Φ}) defined as:

$$\text{epi}_{\Phi} = \{[u \ z]^T : z \geq \Phi(u)\}, \quad (8)$$

where z denotes an upper bound for the function $\Phi(u)$. The projection onto epi_{Φ} is the closest solution to \hat{u} that lies on the boundary of the epigraph set. Since the epigraph set of a convex function is also convex, this projection will yield the global optimum solution. Note that projections onto the epigraph set will yield the solution of the proximal operator only if the search space of the proximal operator is a convex set $\mathcal{U} \in \mathbb{R}^k$ [41]. In practice, a family of solutions can be obtained by introducing a scaling parameter to alter the size of the epigraph set:

$$\text{epi}_{\Phi'} = \{[u \ z]^T : z \geq \beta_{\Phi} \Phi(u)\}. \quad (9)$$

Here, β_{Φ} serves to control the allowed degree of deviation of u^* from \hat{u} . Note that both z^* and u^* are computed via an orthogonal projection of the input onto epi'_{Φ} . Since the scales of z^* and u^* vary proportionately to the scale of \hat{u} , β_{Φ} can be described in absolute terms. $\beta_{\Phi} > 1$ scales down the epigraph set, resulting in a solution u^* that deviates further from \hat{u} . Meanwhile, $0 < \beta_{\Phi} < 1$ scales up epi_{Φ} , resulting in a solution u^* that is closer to \hat{u} , where $u^* = \hat{u}$ as $\beta_{\Phi} \rightarrow 0$. To obtain more conservative solutions, here we used $0 < \beta_{\Phi} < 1$ for both sparsity and TV projections. The resulting projection point determines both the size of the Φ -ball in \mathbb{R}^k (i.e. ℓ_1 -ball or TV-ball, see Fig. 2) and the actual projection onto the ball. Hence, the proximal operator in (6) enables assessing the optimal constraint parameters in (5) using the input vector \hat{u} as explained below.

1) *Self-tuning sparsity projection:* The sparsity projections were implemented using projections onto the epigraph set of the ℓ_1 -norm function, applied on wavelet coefficients. The image coefficients $m_{nd} = \mathcal{F}^{-1}\{x_{nd}\}$ are obtained by inverse Fourier transformation of k-space data, x_{nd} , for acquisition n and coil d . The wavelet coefficients for m_{nd} are then given by $w_{ls,nd} = \Psi_{ls}\{m_{nd}\}$ at subband s and level l , and w_{ls}

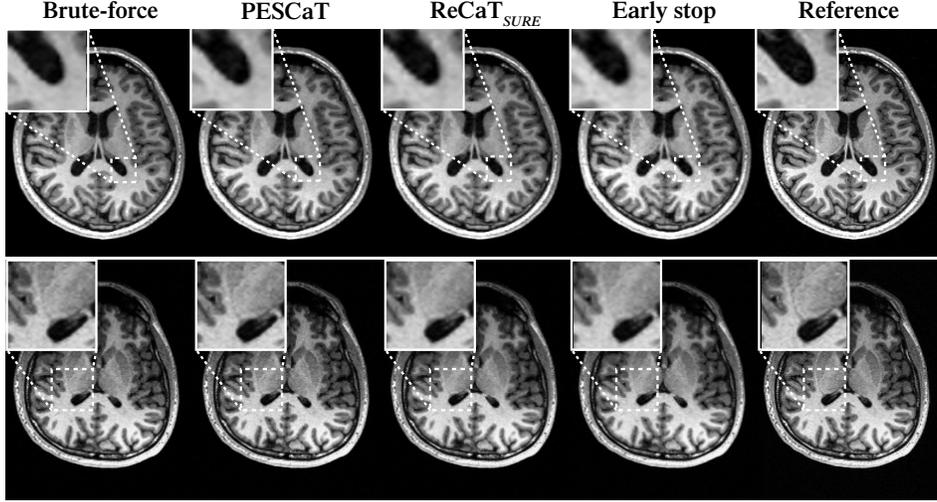


Fig. 5. Reconstructions of in vivo T1-weighted acquisitions of the brain at $R=4$. Brute-force, PESCaT, $\text{ReCaT}_{\text{SURE}}$, $\text{ReCaT}_{\text{SURE}}$ with early stop, and reference images are shown in two representative subjects. White boxes display zoomed-in portions of the images. PESCaT achieves significantly improved image quality compared to $\text{ReCaT}_{\text{SURE}}$ with early stop that was matched PESCaT in terms of the total reconstruction time. Meanwhile, PESCaT yields similar image quality to $\text{ReCaT}_{\text{SURE}}$ and brute-force methods.

denotes the aggregate vector pooling $w_{ls,nd}$ across coils and acquisitions. Assuming $\hat{w} = w_{ls}$ is the input vector, the proximal formulation in (6) becomes:

$$\text{prox}_{\ell_1^2}(\hat{w}) = \arg \min_u \|\hat{w} - u\|_2^2 + \|u\|_1^2. \quad (10)$$

The solution to (10) is then obtained by projecting $[w_{ls} \ 0]^T$ onto the scaled epigraph set (see Fig. 2a):

$$\text{epi}_{\ell_1} = \{[u \ z]^T \in \mathbb{R}^{k+1} : z \geq \beta_{\ell_1} \|u\|_1\}, \quad (11)$$

where β_{ℓ_1} denotes the epigraph scaling factor for the ℓ_1 -norm. As demonstrated in Fig. 2a, the closest orthogonal projection of $[w_{ls} \ 0]^T$ onto the epigraph set lies on the boundary of epi_{ℓ_1} . For the simple case of \mathbb{R}^2 ($k = 1$), $[w_{ls} \ 0]^T$ is projected onto the $z = \beta_{\ell_1} |u|$ line, yielding $z_{ls}^* = \frac{\beta_{\ell_1} |w_{ls}|}{\beta_{\ell_1}^2 + 1}$. It can be shown that for arbitrary k , the z -intercept is:

$$z_{ls}^* = \frac{\beta_{\ell_1} \|w_{ls}\|_1}{\beta_{\ell_1}^2 k + 1}. \quad (12)$$

The value of the z -intercept also determines the size of the respective ℓ_1 -ball, $B_{\ell_1,ls} = \{u \in \mathbb{R}^k : \|u\|_1 \leq \epsilon_{\ell_1,ls}\}$, as:

$$\epsilon_{\ell_1,ls} = \frac{z_{ls}^*}{\beta_{\ell_1}}. \quad (13)$$

Therefore, w_{ls}^* can be computed by finding the projection of w_{ls} onto the ℓ_1 -ball of size $\epsilon_{\ell_1,ls}$. To efficiently implement this projection, we used a soft-thresholding operation [42]:

$$w_{ls}^* = e^{i\angle w_{ls}} \max(|w_{ls}| - \theta_{ls}, 0), \quad (14)$$

where magnitudes of wavelet coefficients are subjected to a threshold of θ_{ls} , and phases of coefficients are individually restored via $e^{i\angle w_{ls}}$. We propose to determine the value of θ_{ls} given $\epsilon_{\ell_1,ls}$ using an efficient ranking algorithm [43]. The proposed algorithm first sorts the absolute values of the wavelet coefficients $w_{ls,nd}$ to attain a rank-ordered sequence $\{\mu_i\}_{i=1}^k$ where $\mu_1 > \mu_2 > \dots > \mu_k$. This sequence is then

analyzed to find the threshold that approximately yields a resultant ℓ_1 -norm of value $\epsilon_{\ell_1,ls}$ in the thresholded coefficients:

$$\rho_{ls} = \max\{j \in \{1, 2, \dots, k\} : \mu_j - \frac{1}{j} \left(\sum_{r=1}^j \mu_r - \epsilon_{\ell_1,ls} \right) > 0\},$$

$$\theta_{ls} = \frac{1}{\rho_{ls}} \left(\sum_{n=1}^{\rho_{ls}} \mu_n - \epsilon_{\ell_1,ls} \right). \quad (15)$$

Note that the determined threshold directly translates to $\lambda_{\ell_1,ls}$ in (4) by [44]:

$$\lambda_{\ell_1,ls} = 2\theta_{ls}. \quad (16)$$

Projections were separately performed for each subband s at each wavelet decomposition level l to determine the respective w_{ls}^* , and $\epsilon_{\ell_1,ls}$. Since wavelet coefficients were pooled across coils and acquisitions, parameter selection is performed jointly across coils and acquisitions. Since the only free parameter in the proposed method is the epigraph scaling constant β_{ℓ_1} , the selection of $3 \times L$ parameters in (5) are transformed into the selection of a single parameter. Here, the optimal β_{ℓ_1} was empirically determined in a group of training subjects and then used to obtain reconstructions in held-out test subjects.

2) *Self-tuning TV projection*: The TV projections were implemented using projections onto the epigraph set of the TV-norm function, applied on image coefficients. Letting $\hat{m} = m_{nd}$ be the input vector, the proximal formulation in (6) becomes:

$$\text{prox}_{\text{TV}^2}(\hat{m}) = \arg \min_u \|\hat{m} - u\|_2^2 + \|u\|_{\text{TV}}^2. \quad (17)$$

The solution to (17) is then obtained by projecting $[m_{nd} \ 0]^T$ onto the scaled epigraph set (see Fig. 2b):

$$\text{epi}_{\text{TV}} = \{[u \ z]^T \in \mathbb{R}^{k+1} : z \geq \beta_{\text{TV}} \|u\|_{\text{TV}}\}, \quad (18)$$

where β_{TV} denotes the epigraph scaling factor for the TV-norm. Unlike the projection onto the ℓ_1 -norm epigraph, projection onto generic epigraph sets (including TV-norm epigraph) does not have a closed-form solution. As demonstrated

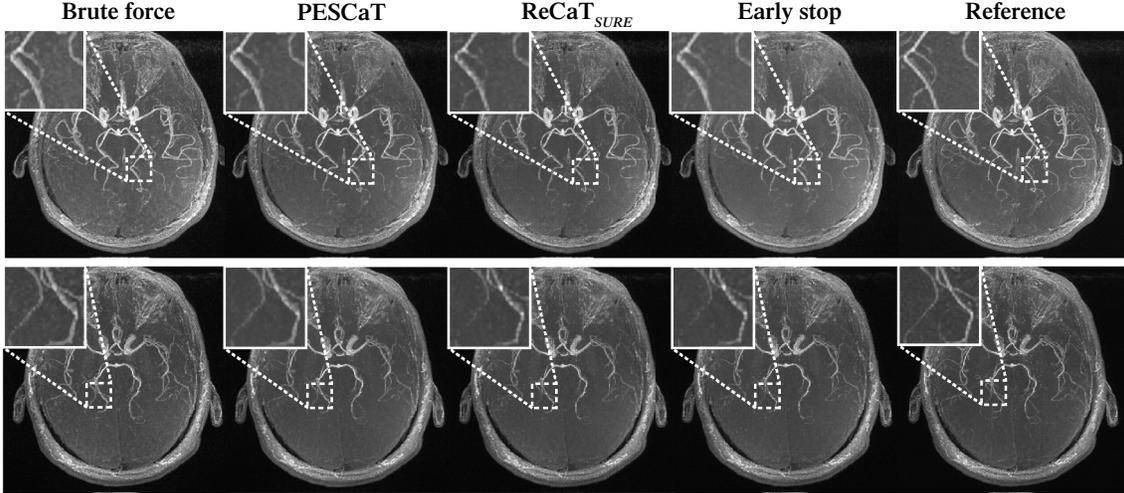


Fig. 6. Reconstructions of in vivo ToF angiography acquisitions of the brain at $R=4$. Maximum-intensity projection (MIP) views of brute-force, PESCaT, $\text{ReCaT}_{\text{SURE}}$, $\text{ReCaT}_{\text{SURE}}$ with early stop, and reference brain volumes are shown in two representative subjects. White boxes display zoomed-in portions of the MIPs. PESCaT yields superior depiction of vasculature compared to both $\text{ReCaT}_{\text{SURE}}$ and $\text{ReCaT}_{\text{SURE}}$ with early stop. It also maintains similar image quality to brute-force reconstructions.

in Fig. 2b, PESCaT uses an iterative epigraphical splitting method to perform the projection efficiently [45]. In the initial step of this approach, complex-valued $m^{(0)} = m_{nd}$ is projected onto the supporting hyperplane of epi_{TV} , at $[m_{nd} \ \beta_{TV} \|m_{nd}\|_{TV}]^T$ resulting in $m^{(1)}$. The supporting hyperplane is determined by evaluating the gradient of the epigraph surface. In the following step, $m^{(1)}$ is projected onto the level set, $L_{TV} = \{[u \ z]^T : z \leq 0\}$, by forcing the last element of $m^{(1)}$ to zero. This projection yields the next estimate $m^{(2)}$. These two projections are iterated. Note that all steps of the splitting procedure are performed in complex domain, thereby, regularizing magnitude and phase channels simultaneously. Previous studies have shown that the second derivative of distance between the input vector and the projections on the supporting hyperplanes ($\|m_{nd} - m^{(2i+1)}\|_2$) is negative as the projections approach to the true projection solution and is positive as the projections deviate from it [41]. Thus, in case of a sign change in the second derivative a refinement step is performed, where $m^{(2i)}$ is projected onto the supporting hyperplane at $\frac{m^{(2i+1)} + m^{(2i-1)}}{2}$. This heuristic approach has been shown to converge to the global solution for TV projections [45]. Note that the projection uniquely specifies the z -intercept, z_{nd}^* . Hence, the size of the corresponding TV-ball, $B_{TV,nd} = \{u \in \mathbb{R}^k : \|u\|_{TV} \leq \epsilon_{TV,nd}\}$, can be calculated as:

$$\epsilon_{TV,nd} = \frac{z_{nd}^*}{\beta_{TV}}. \quad (19)$$

Note that it is nontrivial to explicitly express $\lambda_{TV,nd}$ in (4) in terms of $\epsilon_{TV,nd}$ in (5). Yet, constraining $\epsilon_{TV,nd}$ implicitly enforces a set of regularization parameters $\lambda_{TV,nd}$.

Projections were separately performed for each acquisition n and coil d to determine the respective m_{nd}^* , and $\epsilon_{TV,nd}$. Since the only free parameter is the epigraph scaling constant β_{TV} , the selection of $N \times D$ parameters in (5) is transformed into the selection of a single parameter. Here, the optimal β_{TV} was empirically determined in a group of training subjects and then used to obtain reconstructions in held-out test subjects.

TABLE I
PEAK SIGNAL-TO-NOISE RATIO (PSNR) FOR SIMULATED PHANTOM

	R=2	R=4	R=6
Brute-force	30.29 ± 0.24	28.16 ± 0.27	27.47 ± 0.22
PESCaT	29.56 ± 0.34	26.68 ± 0.24	26.25 ± 0.19
$\text{ReCaT}_{\text{SURE}}$	29.44 ± 0.21	25.52 ± 0.23	24.72 ± 0.16
Early stop	27.67 ± 0.26	24.96 ± 0.23	23.66 ± 0.21

PSNR was measured between the reconstructed image and a fully-sampled reference image. Measurements were obtained for brute-force, PESCaT, $\text{ReCaT}_{\text{SURE}}$ and $\text{ReCaT}_{\text{SURE}}$ with early stop methods. Results are reported as mean \pm std across five cross-sections.

All reconstruction algorithms were executed in MATLAB (MathWorks, MA). The implementations used libraries from the SPIRiT toolbox [28]. The PESCaT algorithm is available for general use at <http://github.com/icon-lab/mrirecon>.

III. METHODS

A. Alternative reconstructions

To demonstrate the performance of PESCaT, we compared it against several alternative reconstructions that aim to select regularization parameters.

a) Self-tuning regularized ReCaT ($\text{ReCaT}_{\text{SURE}}$): Our previously proposed multi-coil multi-acquisition method (ReCaT) did not include any regularization parameters [33]. We have implemented a variant of ReCaT incorporating sparsity and TV regularization terms where the regularization parameters are automatically selected using the data. This reconstruction method iteratively synthesizes unacquired data as weighted combinations of collected data across coils and/or acquisitions. It uses sparsity and TV projections to enforce image sparsity. At each iteration, the regularization parameter for the sparsity term is selected based on the SURE criterion. The regularization parameter for the TV term is selected based on the local standard deviation of the reconstructed image from the previous iteration.

An alternating projections onto sets algorithm was used in $\text{ReCaT}_{\text{SURE}}$ to solve the reconstruction problem cast in (4). $\text{ReCaT}_{\text{SURE}}$ used a single sparsity regularization parameter for all subbands and levels of wavelet coefficients. The sparsity

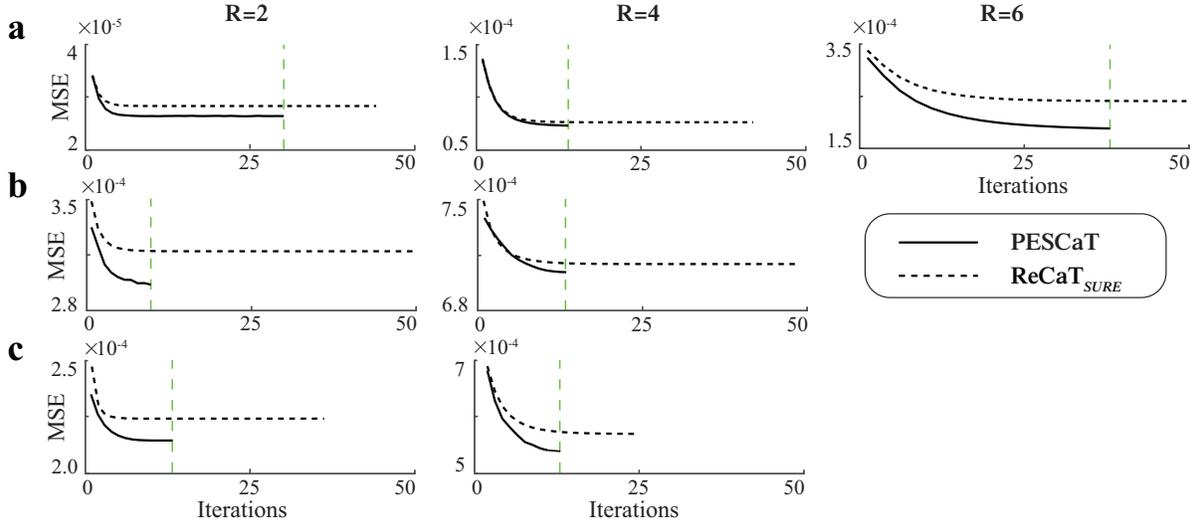


Fig. 7. Convergence behavior of self-tuning reconstructions was assessed on in vivo (a) bSSFP, (b) T1-weighted, and (c) ToF acquisitions of the brain. Mean-squared error (MSE) was calculated between the image reconstructed at each iteration and the fully-sampled reference image. The progression of MSE across iterations is shown for a representative cross-section reconstructed using PESCaT (solid lines) and ReCaT_{SURE} (dashed lines) at R=2 (left), 4 (middle), and 6 (right). Reconstructions were stopped once convergence criteria were reached (see *Methods*). The iterations at which PESCaT converged are indicated (dashed green lines). In all cases, PESCaT converges in a significantly smaller number of iterations, and it converges to a solution with lower MSE than ReCaT_{SURE}.

regularization parameter was determined via a line search over the range $[2 \times 10^{-5}, 2 \times 10^{-1}]$. The TV regularization parameter was taken as one-third of the median local standard deviation [46]. All remaining reconstruction parameters were kept identical to PESCaT.

b) *ReCaT_{SURE} with early stop*: The projections performed in each iteration of PESCaT do not involve any line searches, and therefore they are more efficient compared to ReCaT_{SURE}. To enable a fair comparison, we implemented a variant of ReCaT_{SURE} that was stopped once the total reconstruction time reached that of PESCaT. All reconstruction parameters except the total number of iterations were kept identical to ReCaT_{SURE}.

c) *ReCaT with empirically-tuned parameters (ReCaT_{fixed})*: To demonstrate the effects of prescribing separate regularization parameters for different subbands/levels or coils/acquisitions in PESCaT, we implemented a variant of ReCaT with a single sparsity parameter across all subbands/levels and a single TV parameter across all coils/acquisitions. Similar to PESCaT, this reconstruction method iteratively synthesizes unacquired data as weighted combinations of collected data across coils and/or acquisitions. ReCaT_{fixed} was tuned using held-out data. The sparsity and TV parameters were independently varied across a broad range $[10^{-5}, 0.5]$. Separate reconstructions were obtained for each parameter set, and reconstruction quality was taken as peak signal-to-noise ratio (PSNR) between the reconstructed image and the fully-sampled reference image. The parameter set that yielded the maximum PSNR was selected. Sparsity and TV parameters were fixed across iterations. All remaining reconstruction parameters were kept identical to PESCaT.

d) *Brute-force reconstruction*: To evaluate the success of PESCaT in selecting the optimal parameters, a brute-force reconstruction was implemented to solve the problem in (4). The brute-force method used a constant set of regularization

parameters across iterations. The sparsity and TV parameters were independently varied across the range $[10^{-5}, 0.5]$. Separate reconstructions were obtained for each parameter set, and reconstruction quality was taken as PSNR between the reconstructed image and the fully-sampled reference image. The parameter set that yielded the maximum PSNR was selected. All remaining reconstruction parameters were kept identical to PESCaT.

e) *ESPIRiT with PES parameter tuning (PESSPIRiT)*:

To compare the performance of PESCaT against conventional parallel imaging, we implemented a variant of ESPIRiT [47] that included sparsity and TV regularization terms. Similar to ESPIRiT, this method iteratively reconstructs images based on coil sensitivities estimated from central calibration data. In each iteration, the sparsity and TV regularization parameters were tuned using PES. Two other variants, PESSPIRiT with only the sparsity regularization (PESSPIRiT_{ℓ₁}) and PESSPIRiT with only the TV regularization (PESSPIRiT_{TV}) were also implemented. In all variants, the stopping criterion was identical to PESCaT to enable a fair comparison.

B. Simulations

Simulations were performed using a realistic brain phantom at 0.5 mm isotropic resolution (<http://www.bic.mni.mcgill.ca/brainweb>). Phase-cycled bSSFP signals were assumed with T1/T2: 3000/1000 ms for cerebrospinal fluid, 1200/250 ms for blood, 1000/80 ms for white matter, 1300/110 ms for gray matter, 1400/30 ms for muscle, and 370/130 ms for fat [30]. Single-coil three-dimensional (3D) acquisitions were assumed with TR/TE=5.0/2.5 ms, flip angle=45°, and phase-cycling increments $\Delta\phi=2\pi \frac{[0:1:N-1]}{N}$. We used a simulated field inhomogeneity distribution corresponding to an off-resonance shift with zero mean and 62 Hz standard deviation. A bivariate Gaussian noise was added to simulated acquisitions to attain signal-to-noise ratio (SNR)=20, where SNR was taken as the ratio of the mean power in the phantom image to the

TABLE II
PSNR AND NRMSE FOR IN VIVO BSSFP DATASET

	R=2		R=4		R=6	
	PSNR	NRMSE $\times 10^3$	PSNR	NRMSE $\times 10^3$	PSNR	NRMSE $\times 10^3$
Brute-force	44.31 \pm 0.72	7.31 \pm 0.23	40.21 \pm 0.78	11.18 \pm 0.45	37.62 \pm 0.61	16.79 \pm 0.56
PESCaT	43.93 \pm 0.65	8.16 \pm 0.27	39.64\pm0.61	11.93\pm0.43	36.72\pm0.29	18.36\pm0.34
ReCaT_{fixed}	44.11\pm0.62	7.57\pm0.33	39.08 \pm 0.47	12.59 \pm 0.48	36.51 \pm 0.48	18.45 \pm 0.95
ReCaT_{SURE}	42.20 \pm 0.78	10.09 \pm 0.59	37.89 \pm 0.76	16.19 \pm 0.79	35.15 \pm 0.59	20.44 \pm 0.82
Early stop	41.83 \pm 0.66	10.37 \pm 0.52	36.82 \pm 0.63	18.18 \pm 0.77	34.01 \pm 0.32	22.08 \pm 0.31
PESSPIRiT	43.37 \pm 0.44	8.80 \pm 0.35	38.39 \pm 0.59	14.37 \pm 0.84	35.06 \pm 0.35	20.56 \pm 0.71
PESSPIRiT_{ℓ_1}	41.11 \pm 0.62	10.84 \pm 0.62	35.96 \pm 0.47	19.36 \pm 0.86	33.15 \pm 0.55	23.05 \pm 1.40
PESSPIRiT_{TV}	42.41 \pm 0.67	9.61 \pm 0.59	35.57 \pm 0.68	19.88 \pm 1.35	32.24 \pm 0.61	24.28 \pm 1.69

PSNR and NRMSE were measured between the reconstructed image and a fully-sampled reference image. Measurements were obtained for brute-force, PESCaT, ReCaT_{SURE}, ReCaT_{SURE} with early stop, ReCaT_{fixed}, and variants of PESSPIRiT methods. Results are averaged across three subjects, and reported as mean \pm std across five cross-sections.

noise variance. Data were undersampled by a factor (R) of 2, 4, and 6 in the two phase-encode directions using disjoint, variable density random undersampling [4] and normalized so that zero-filled density compensated k-space data had unity norm [1]. Reconstruction quality was taken as PSNR between reconstructions and a fully-sampled reference. To prevent bias, the 98th percentile of image intensities were adjusted to [0, 1]. PSNR values were then averaged across five central axial cross-sections.

To examine the effect of noise on optimal regularization parameters, we performed experiments on the simulated brain phantom where the noise level was systematically varied. The simulations output single-coil single-acquisition brain images with SNR varying in the range [5, 25]. Data were undersampled by R=2, 4, and 6 in the two phase-encode directions using disjoint, variable density random undersampling. Multiple separate reconstructions were obtained for each undersampled dataset via ReCaT_{fixed}, while ℓ_1 and TV regularization parameters were independently varied in the range [0.001, 0.1]. At each SNR level, fully-sampled data were used as reference. PSNR was measured between the reconstructions and the reference. The optimal regularization parameters were selected according to PSNR.

To examine the reliability of the epigraph scaling parameters against noise, reconstructions of the brain phantom were obtained at three separate levels of SNR = 10, 18, 25. Meanwhile, β_{ℓ_1} was varied in the range [0.05, 0.6] and β_{TV} was varied in the range [0.1, 1]. To examine the reliability of the epigraph scaling parameters against variations in the level of detail and spatial resolution, we performed experiments on a simulated numerical phantom. A circular phantom of radius 125 voxels (for a 256 \times 256 field of view) was designed with the background resembling muscle tissue and vertical bright bars of width 12 and height [190, 220, 238, 238, 220, 190] voxels resembling blood vessels (Supp. Fig. 1)¹. Phase-cycled bSSFP signals were assumed with T1/T2: 870/47 ms for muscle, and 1273/259 ms for blood. Three dimensional acquisitions were assumed with TR/TE=4.6/2.3 ms, flip angle=60 $^\circ$, and phase-cycling increments $\Delta\phi=2\pi\frac{[0:1:N-1]}{N}$. A simulated field inhomogeneity distribution corresponding to an off-resonance shift with zero mean and 62 Hz standard deviation was used. Level of detail was varied from low to

high by incrementally placing [1, 3, 6] vertical bars in the phantom. Spatial resolution was varied from low to high by low-pass filtering k-space data to select circular regions of radius [20, 55, 125] voxels. Reconstructions were obtained while β_{ℓ_1} and β_{TV} were varied in the range [0.05, 0.5].

C. In vivo experiments

Experiments were performed to acquire 3D multi-coil multi-acquisition phase-cycled bSSFP, and multi-coil single-acquisition T1-weighted and time-of-flight (ToF) angiography data in the brain. Data were collected on a 3T Siemens Magnetom scanner (maximum gradient strength of 45 mT/m and slew rate of 200 T/m/s). bSSFP and ToF data were collected using a 12-channel receive-only head coil that was hardware compressed to 4 channels. T1-weighted data were collected using a 12-channel receive-only head coil. Separate bSSFP datasets were also collected using a 32-channel head coil. Balanced SSFP data were acquired using a bSSFP sequence with the following parameters: flip angle=30 $^\circ$, TR/TE=8.08/4.04 ms, field-of-view (FOV)=218 mm \times 218 mm, matrix size of 256 \times 256 \times 96, resolution of 0.9 mm \times 0.9 mm \times 0.8 mm, right/left readout direction, and N=8 separate acquisitions with phase-cycling values in the range [0, 2 π) in equispaced intervals. Total acquisition time for the bSSFP sequence was 20:56. T1-weighted data were acquired using an MP-RAGE sequence with the parameters: flip angle=9 $^\circ$, TR/TE=2300/2.98 ms, TI=900 ms, FOV= 256 mm \times 240 mm, matrix size of 256 \times 240 \times 160, resolution of 1.0 mm \times 1.0 mm \times 1.2 mm, and superior/inferior readout direction. Total acquisition time for the MP-RAGE sequence was 9:14. ToF angiograms were acquired using a multiple overlapping thin-slab acquisition (MOTSA) sequence with parameters: flip angle=18 $^\circ$, TR/TE=38/3.19 ms, FOV=204 mm \times 204 mm, matrix size of 256 \times 256 \times 75, isotropic resolution of 0.8 mm, and anterior/posterior readout direction. Total acquisition time for the MOTSA sequence was 14:16. The imaging protocols were approved by the local ethics committee, and all six participants gave written informed consent.

Phase-cycled bSSFP acquisitions with 4 channels were retrospectively undersampled at R=2, 4, and 6. Following phase-cycles were selected: $\Delta\phi = 2\pi\frac{[0:1:N-1]}{N}$ for N=2 and 4, and [0, $\frac{\pi}{2}$, $\frac{3\pi}{4}$, π , $\frac{5\pi}{4}$, $\frac{7\pi}{4}$] for N=6. For this bSSFP dataset, N=R was used. T1-weighted and ToF acquisitions were retrospectively undersampled at R=2 and 4 (note that in these cases N=1). Undersampling was performed across the two phase

¹supplementary materials are available in the supplementary files /multi-media tab.

TABLE III
PSNR AND NRMSE FOR IN VIVO T1-WEIGHTED DATASET

	R=2		R=4	
	PSNR	NRMSE $\times 10^3$	PSNR	NRMSE $\times 10^3$
Brute-force	36.75 \pm 0.55	18.25 \pm 0.89	32.15 \pm 0.41	23.75 \pm 0.85
PESCaT	35.62\pm0.95	19.86\pm0.99	31.44\pm1.09	27.75\pm1.61
ReCaT_{fixed}	35.27 \pm 0.67	20.02 \pm 0.85	30.89 \pm 0.63	29.59 \pm 1.31
ReCaT_{SURE}	35.02 \pm 0.93	20.64 \pm 0.94	30.67 \pm 0.93	30.10 \pm 1.57
Early stop	34.64 \pm 1.03	21.47 \pm 0.95	30.03 \pm 1.06	30.81 \pm 1.65
PESSPIRiT	35.21 \pm 0.74	20.14 \pm 1.10	29.92 \pm 0.63	31.66 \pm 1.60
PESSPIRiT$_{\ell_1}$	31.75 \pm 0.66	27.37 \pm 1.42	27.39 \pm 0.80	34.37 \pm 2.72
PESSPIRiT_{TV}	34.93 \pm 0.69	21.04 \pm 1.03	29.48 \pm 0.60	31.83 \pm 1.61

PSNR and NRMSE were measured between the reconstructed image and a fully-sampled reference image. Measurements were obtained for brute-force, PESCaT, ReCaT_{SURE}, ReCaT_{SURE} with early stop, ReCaT_{fixed}, and variants of PESSPIRiT methods. Results are averaged across three subjects, and reported as mean \pm std across five cross-sections.

encode directions: superior/inferior and anterior/posterior for bSSFP, right/left and anterior/posterior for T1-weighted, superior/inferior and right/left for ToF. Data were normalized so that zero-filled density compensated k-space data had unity norm. Entire volumes were reconstructed, five axial cross-sections equispaced across the entire brain were selected for quantitative assessment. PSNR and normalized root mean-squared error (NRMSE) measurements were averaged across cross-sections.

To investigate the convergence behavior of PESCaT, we studied the evolution of the three cost terms in (4) separately (Supp. Fig. 2). Normalized cost terms associated with calibration consistency, sparsity, and TV terms at the end of each iteration were plotted across iterations. In all datasets, all cost terms diminish smoothly.

To optimize epigraph scaling constants for ℓ_1 - and TV-norm functions, PESCaT was performed on data acquired from three subjects reserved for this purpose. Volumetric reconstructions were performed at R=2, 4, and 6 for bSSFP datasets, and R=2 and 4 for T1-weighted and ToF datasets. Separate reconstructions were obtained while β_{ℓ_1} was varied in the range [0.1, 1], and β_{TV} was varied in the range [0.05, 0.6]. PSNR was measured between the reconstructed and fully-sampled reference images (Supp. Figs. 3, 4). Consistently across subjects and different types of datasets, PSNR values within 95% of the optimum value were maintained in the range $\beta_{\ell_1} = [0.1, 0.3]$, and $\beta_{TV} = [0.2, 0.4]$. Near-optimal PSNR values were attained around $\beta_{\ell_1} = 0.2$ and $\beta_{TV} = 0.3$. Thus, these scaling constants were prescribed for reconstructions thereafter.

To demonstrate the reconstruction performance of PESCaT at high acceleration rates, phase-cycled bSSFP acquisitions with 32 channels were analyzed. This bSSFP dataset was retrospectively undersampled at R= 8, 10 (where N=8). Entire volumes were reconstructed, and PSNR and NRMSE measurements were averaged across five axial cross-sections.

IV. RESULTS

A. Simulations

MRI data may show differential noise and structural characteristics for separate coils and acquisitions, or for separate wavelet subbands and levels. In turn, the optimal regularization parameters can vary across each of these dimensions. To test this prediction, we performed experiments on the simulated

brain phantom, where the noise level was systematically varied and ReCaT_{fixed} reconstructions were performed. For both ℓ_1 and TV regularization, the optimal regularization parameters show a clear increasing trend as SNR is lowered (Supp. Fig. 5). These results suggest that prescribing a fixed parameter can cause performance loss when a good compromise cannot be achieved across subbands/levels or coils/acquisitions. It can also render the reconstruction more susceptible to deviations from the optimal value of the regularization parameter.

In contrast, PESCaT uses only two global parameters to control the overall sparsity of the solutions in wavelet domain (β_{ℓ_1}) and TV domain (β_{TV}). Given these scaling parameters, regularization parameters for individual subbands/levels and coils/acquisitions are determined adaptively in a data-driven manner. To examine the reliability of the scaling parameters against noise, reconstructions were obtained at varying SNR levels. The PSNR curves as a function of β_{ℓ_1} and β_{TV} demonstrate substantial flatness, yielding near-optimal performance across the entire range of values examined (Supp. Fig. 6). To further examine the reliability of the scaling parameters against variations in the level of detail and spatial resolution, reconstructions were obtained at low, medium and high levels of detail and resolution. Supp. Fig. 7 displays PSNR across β_{ℓ_1} and Supp. Fig. 8 displays PSNR across β_{TV} values. Again, PSNR curves as a function of β_{ℓ_1} and β_{TV} demonstrate substantial flatness, yielding near-optimal performance across the entire range of values examined.

Following these basic demonstrations, PESCaT was performed on bSSFP acquisitions of a simulated brain phantom. Representative reconstructions and error maps for PESCaT and ReCaT_{SURE} with R=6 are shown in Fig. 3. PESCaT yields reduced error across the FOV compared to ReCaT_{SURE}. This improvement with PESCaT becomes further noticeable when ReCaT_{SURE} is stopped early to match its reconstruction time to PESCaT. Quantitative assessments of image quality at R=2, 4, and 6 are listed in Table I. Among all techniques tested, PESCaT achieves the most similar performance to the time-consuming brute-force reconstruction. On average, PESCaT improves PSNR by 0.87 ± 0.74 dB over ReCaT_{SURE} and by 1.87 ± 0.73 dB over ReCaT_{SURE} with early stop (mean \pm std. across five cross-sections, average of R=2, 4, 6). Note that the proposed method attains near-optimal performance while enabling improved computational efficiency. The average reconstruction time per slice is 27 ± 5 s for ReCaT_{SURE} and only

TABLE IV
PSNR AND NRMSE FOR IN VIVO TOF DATASET

	R=2		R=4	
	PSNR	NRMSE $\times 10^3$	PSNR	NRMSE $\times 10^3$
Brute-force	36.57 ± 1.61	18.61 ± 0.71	33.08 ± 1.55	22.75 ± 0.97
PESCaT	36.32 ± 1.14	18.79 ± 0.57	31.86 ± 1.21	27.55 ± 1.03
ReCaT_{fixed}	35.86 ± 0.60	19.41 ± 0.86	30.51 ± 0.60	30.17 ± 1.60
ReCaT_{SURE}	35.55 ± 1.25	19.97 ± 0.94	31.34 ± 1.13	28.40 ± 1.57
Early stop	35.46 ± 1.12	20.14 ± 0.95	27.69 ± 1.19	28.71 ± 1.65
PESSPIRiT	35.86 ± 0.57	19.42 ± 0.90	30.83 ± 0.47	29.72 ± 1.41
PESSPIRiT_{ℓ_1}	32.00 ± 0.55	24.89 ± 1.46	27.22 ± 0.58	37.84 ± 2.71
PESSPIRiT_{TV}	35.66 ± 0.53	19.72 ± 0.84	30.41 ± 0.64	30.42 ± 1.89

PSNR and NRMSE were measured between the reconstructed image and a fully-sampled reference image. Measurements were obtained for brute-force, PESCaT, ReCaT_{SURE}, ReCaT_{SURE} with early stop, ReCaT_{fixed}, and variants of PESSPIRiT methods. Results are averaged across three subjects, and reported as mean \pm std across five cross-sections.

7 ± 4 s for PESCaT, resulting in a 4-fold gain in efficiency for the phantom dataset.

B. In vivo experiments

We first examined the evolution of the cost terms during PESCaT reconstruction of in vivo bSSFP and T1-weighted datasets (Supp. Figs. 9, 10). Both ℓ_1 and TV cost terms decrease towards later iterations indicating that the images better conform to a compressible representation.

Next, PESCaT was demonstrated for in vivo bSSFP, T1-weighted, and ToF imaging of the brain. Representative reconstructions with R=6 for bSSFP and R=4 for T1-weighted and ToF acquisitions are displayed in Figs. 4, 5, and 6. Representative reconstructions of individual phase cycles in the bSSFP dataset, and of cross-sections in the ToF dataset are shown in Supp. Fig. 11. Overall, PESCaT and ReCaT_{SURE} reconstructions perform similar to the brute-force optimized reconstructions. Yet, PESCaT yields slightly lower levels of residual aliasing in comparison to ReCaT_{SURE}, and this difference is particularly noticeable for visualization of small vessels in ToF images (Fig. 6). The improvement in reconstruction quality with PESCaT is more prominent when ReCaT_{SURE} is stopped early to match its reconstruction time to PESCaT.

Quantitative assessments of the in vivo reconstructions are listed in Tables II, III, and IV. For all datasets and R, PESCaT yields the closest performance to the brute-force reconstruction among alternative self-tuning methods. For bSSFP datasets, PESCaT improves PSNR by 1.23 ± 0.29 dB over ReCaT_{SURE} and by 2.55 ± 0.51 dB over ReCaT_{SURE} with early stop (mean \pm std. across three subjects, average of R=2, 4, 6). For T1-weighted datasets, PESCaT improves PSNR by 0.71 ± 0.25 dB over ReCaT_{SURE} and by 1.21 ± 0.43 dB over ReCaT_{SURE} with early stop (mean \pm std. across three subjects, average of R=2, 4). For ToF datasets, PESCaT improves PSNR by 0.72 ± 0.46 dB over ReCaT_{SURE} and by 0.94 ± 0.51 dB over ReCaT_{SURE} with early stop (mean \pm std. across three subjects, average of R=2, 4). Compared to empirically-tuned ReCaT_{fixed}, PESCaT improves PSNR by 0.20 ± 0.37 dB for bSSFP datasets, by 0.45 ± 0.14 dB for T1-weighted datasets, and by 0.91 ± 0.63 dB for ToF datasets (mean \pm std. across three subjects, average of R=2, 4, 6 for bSSFP, average of R=2, 4 for T1-weighted and ToF datasets). Because both methods were allowed to optimize parameters in training subjects, these results suggest that selecting different regularization pa-

rameters for each coil/acquisition/subband/level improves reconstruction performance. Compared to PESSPIRiT, PESCaT improves PSNR by 1.16 ± 0.55 dB for bSSFP datasets, by 0.97 ± 0.78 dB for T1-weighted datasets, and by 0.76 ± 0.40 dB for ToF datasets (mean \pm std. across three subjects, average of R=2, 4, 6 for bSSFP, average of R=2, 4 for T1-weighted and ToF datasets). Performance enhancement is even more prominent compared to PESSPIRiT variants that only include sparsity or TV regularization.

To assess the computational efficiency of self-tuning methods, representative reconstructions were performed for a single cross-section of in vivo bSSFP, T1-weighted, and ToF acquisitions. The true MSE between the reconstructed and fully-sampled reference images were recorded across iterations of PESCaT and ReCaT_{SURE}. MSE curves across iterations are displayed in Fig. 7. Compared to ReCaT_{SURE}, the proposed method converges to a lower MSE value for all R and datasets. Furthermore, PESCaT reduces the number of iterations by 43.3% for bSSFP (average over R=2, 4, 6), 74.5% for T1-weighted (average over R=2, 4) and 53.2% for ToF (average over R=2, 4) datasets. Note that each iteration of PESCaT performs more efficient geometric projections without explicit parameter searches. The reconstruction times for PESCaT and alternative methods are listed in Supp. Table I. On average, the reconstruction time of ReCaT_{SURE} was 1641 ± 45 s for bSSFP, 1799 ± 66 s for T1-weighted, and 565 ± 58 s for ToF datasets (mean \pm std. across five cross-sections, average over R=2, 4 for T1-weighted and ToF imaging; R=2, 4, 6 for bSSFP imaging). In contrast, the reconstruction time of PESCaT was merely 164 ± 25 s for bSSFP, 196 ± 44 s for T1-weighted, and 159 ± 32 s for ToF datasets. These results suggest that PESCaT offers up to 10-fold gain in efficiency compared to the alternative self-tuning method ReCaT_{SURE}. While PESSPIRiT yields similar reconstruction times and ReCaT_{fixed} slightly reduces reconstruction times compared to PESCaT, both methods yield inferior reconstruction quality.

Lastly, reconstruction performance of PESCaT was demonstrated at higher acceleration rates using the 32-channel bSSFP datasets (Supp. Fig. 12 and Supp. Table II). The proposed method improves PSNR by 0.14 ± 0.04 compared to ReCaT_{fixed}, by 1.59 ± 0.45 compared to ReCaT_{SURE}, by 3.77 ± 0.61 compared to ReCaT_{SURE} with early stop, and by 4.08 ± 0.55 over PESSPIRiT (mean \pm std. across three subjects, average of R=8, 10). These results help demonstrate the utility

of PESCiT in enabling higher acceleration factors when using modern coil arrays.

V. DISCUSSION

In this study, we have proposed a new self-tuning method for CS reconstruction of single-coil multi-acquisition, multi-coil single-acquisition, and multi-coil multi-acquisition datasets. The proposed method performs sparsity projections across coils and acquisitions to penalize the ℓ_1 -norm of wavelet coefficients, and TV projections to penalize the finite-differences gradients of image coefficients. Separate sparsity regularization parameters are selected at each wavelet subband and level, and separate TV regularization parameters are selected at each coil and acquisition. Efficient projections onto the boundary of the epigraph sets of the ℓ_1 -norm and TV-norm functions are used to simultaneously calculate the projections themselves and automatically determine the relevant regularization parameters. PESCiT does not have any constraints regarding the number of acquisitions or coils. Therefore, it can be readily applied to both single-acquisition and multi-acquisition datasets regardless of the number of coils available. PESCiT also offers flexibility regarding the inclusion of regularization terms. Because the algorithm has a modular structure with respect to individual calibration, sparsity, and TV projections, it is possible to omit either TV or sparsity regularization. The proposed method will still work towards a solution at the intersection of the remaining projection sets.

In a recent study, we proposed a reconstruction for multi-coil multi-acquisition bSSFP imaging, named ReCaT [46]. Here, we have implemented a self-tuning version of ReCaT (ReCaT_{SURE}). Similar to PESCiT, ReCaT_{SURE} uses sparsity projections implemented via soft-thresholding and TV projections implemented via iterative clipping. However, in ReCaT_{SURE}, the sparsity regularization parameter was selected via a SURE-based method to minimize the expected reconstruction error. TV regularization parameter was selected in a data-driven manner based on the local standard deviations within the reconstructed image. Since parameter selection in ReCaT_{SURE} involves line searches over a relevant range of parameters, it can be computationally expensive. In contrast, PESCiT leverages highly efficient geometric projections onto epigraph sets to simultaneously select the optimal parameters and calculate the projections. Hence, PESCiT enables significant savings in reconstruction time compared to self-tuning methods based on line searches. Meanwhile, the main advantage of PESCiT over an empirically-tuned reconstruction that optimizes regularization parameters on training data is that it allows for independent selection of regularization parameters for each coil/acquisition/subband/level. The superior reconstruction quality of PESCiT compared to ReCaT_{fixed} confirms this prediction quantitatively.

The proposed method includes two epigraph scaling constants β_{ℓ_1} and β_{TV} as free parameters. Here we have empirically demonstrated that the optimal scaling constants are highly consistent across individual subjects, across different noise levels and across multiple imaging contrasts of the same anatomy. These observations are also complemented by prior work that suggests that the solutions of epigraph

sets projections are robust against deviations from optimal scaling constants [45], [48]. It remains to be demonstrated whether the scaling constants are also similar across different anatomies. Still, we expect that PESCiT shows improved robustness against variability in datasets compared to the empirically-tuned ReCaT_{fixed}. The optimal regularization parameters for ReCaT_{fixed} showed relatively high variability across the datasets examined in this study (not shown). Thus, ReCaT_{fixed} might require more careful tuning of regularization parameters, resulting in relatively higher computational overhead.

Further performance improvements might be attained by addressing some limitations of the proposed method. For multi-acquisition datasets, significant motion among acquisitions can reduce reconstruction quality. A motion-correction projection can be incorporated into the PESCiT algorithm to mitigate artifacts due to the residual motion. Second, the proposed method uses a fully-sampled central region in k-space to estimate the tensor interpolation kernel. In applications where the acquisition of calibration data is impractical such as spectroscopic and dynamic imaging, calibrationless approaches could be incorporated for improved performance [49], [50]. Third, although the epigraph scaling constants β_{ℓ_1} and β_{TV} were optimized over a held-out dataset, it might be possible to automatically select them using parameter selection via SURE or GCV. This remains an important future research direction toward fully-automated reconstructions.

Here, the alternating projections onto sets algorithm was used to find a solution at the intersection point of the sets corresponding to calibration, sparsity, and TV projections. Rapid convergence was observed in all examined cases. However, in situations where the intersection between these sets is sparsely populated, more sophisticated algorithms such as alternating direction method of multipliers (ADMM) could be used for fast and effective optimization [51]. PESCiT employs projections onto epigraph sets to concurrently select regularization parameters and perform projections. As such, it is non-trivial to efficiently adapt the proposed parameter selection to an ADMM-based reconstruction. It remains an important future work to benchmark PESCiT against ADMM coupled with an appropriate parameter-selection strategy.

Projections onto epigraph sets were used to penalize ℓ_1 -norm and TV-norm functions in this study. Note that the projection onto convex sets formulation allows penalization of any convex function. Thus, the proposed technique could be generalized to include alternative regularizers such as filtered variation or total generalized variation [52], [53]. These modifications might allow performance enhancements in applications where standard TV regularization yields undesirable block artifacts.

In conclusion, PESCiT enables near-optimal image quality while automatically selecting regularization parameters in reconstructions of undersampled MRI datasets. Parameter selection for ℓ_1 -norm and TV-norm regularizers and projections onto the ℓ_1 and TV-balls are performed simultaneously. PESCiT was demonstrated to outperform alternative self-tuning approaches based on SURE in bSSFP, T1-weighted and time-of-flight angiographic imaging. The results presented

here demonstrate that PESCiT is a promising method for CS-MRI in routine practice.

REFERENCES

- [1] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: the application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–95, 2007.
- [2] K. Block, M. Uecker, and J. Frahm, "Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint," *Magnetic Resonance in Medicine*, vol. 57, no. 6, pp. 1086–1098, 2007.
- [3] B. Liu, Y. Zou, and L. Ying, "SparseSENSE: application of compressed sensing in parallel MRI," *Proceedings of the International Conference on Information Technology and Application in Biomedicine*, pp. 127–130, 2008.
- [4] T. Çukur, "Accelerated phase-cycled SSFP imaging with compressed sensing," *IEEE Transactions on Medical Imaging*, vol. 34, no. 1, pp. 107–115, 2015.
- [5] P. Hansen, "Analysis of discrete ill-posed problems by means of the L-curve," *SIAM Review*, vol. 34, no. 4, pp. 531–580, 1992.
- [6] P. C. Hansen and D. P. O’Leary, "The use of the L-curve in the regularization of discrete ill-posed problems," *SIAM Journal on Scientific Computing*, vol. 14, no. 6, pp. 1487–1503, 1993.
- [7] T. Reginska, "A regularization parameter in discrete ill-posed problems," *SIAM Journal on Scientific Computing*, vol. 17, no. 3, pp. 740–749, 1996.
- [8] X. Zhang and M. Desal, "Adaptive denoising based on SURE risk," *IEEE Signal Processing Letters*, vol. 5, no. 10, pp. 265–267, 1998.
- [9] F. Lin, K. K. Kwong, J. W. Belliveau, and L. L. Wald, "Parallel imaging reconstruction using automatic regularization," *Magnetic Resonance in Medicine*, vol. 51, no. 3, pp. 559–567, 2004.
- [10] L. Ying, D. Xu, and Z. Liang, "On Tikhonov regularization for image reconstruction in parallel MRI," *Proceedings of the 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, vol. 1, pp. 1056–1059, 2004.
- [11] B. Bilgic, I. Chatnuntawech, and C. Langkammer, "Sparse methods for quantitative susceptibility mapping," *Proceedings of the SPIE*, vol. 9597, pp. 711–722, 2015.
- [12] B. Bilgic, I. Chatnuntawech, A. P. Fan, K. Setsompop, S. F. Cauley, L. L. Wald, and E. Adalsteinsson, "Fast image reconstruction with l_2 -regularization," *Magnetic Resonance Imaging*, vol. 40, no. 1, pp. 181–191, 2014.
- [13] C. Vogel, "Non-convergence of the L-curve regularization parameter selection method," *Inverse Problems*, vol. 12, no. 5, p. 535, 1996.
- [14] B. Kressler, D. L. Rochefort, and T. Liu, "Nonlinear regularization for per voxel estimation of magnetic susceptibility distributions from MRI field maps," *IEEE Transactions on Medical Imaging*, vol. 29, no. 2, pp. 273–281, 2010.
- [15] R. Giryes, M. Elad, and Y. Eldar, "The projected GSURE for automatic parameter tuning in iterative shrinkage methods," *Applied and Computational Harmonic Analysis*, vol. 30, no. 3, pp. 407–422, 2011.
- [16] G. Golub, M. Heath, and G. Wahba, "Generalized cross-validation as a method for choosing a good ridge parameter," *Technometrics*, vol. 21, no. 2, pp. 215–223, 1979.
- [17] C. Stein, "Estimation of the mean of a multivariate normal distribution," *The Annals of Statistics*, vol. 9, no. 6, pp. 1135–1151, 1981.
- [18] J. D. Carew, G. Wahba, X. Xie, E. V. Nordheim, and M. E. Meyerand, "Optimal spline smoothing of fMRI time series by generalized cross-validation," *NeuroImage*, vol. 18, pp. 950–961, 2003.
- [19] S. Sourbron, R. Luypaert, and V. Schuerbeek, "Choice of the regularization parameter for perfusion quantification with MRI," *Physics in Medicine and Biology*, vol. 49, no. 14, pp. 3307–3324, 2004.
- [20] S. Sourbron, R. Luypaert, P. Van Schuerbeek, M. Dujardin, T. Stadnik, and M. Osteaux, "Deconvolution of dynamic contrast-enhanced MRI data by linear inversion: Choice of the regularization parameter," *Magnetic Resonance in Medicine*, vol. 52, no. 1, pp. 209–213, 2004.
- [21] S. Ramani, Z. Liu, J. Rosen, J. Nielsen, and J. A. Fessler, "Regularization parameter selection for nonlinear image restoration and MRI reconstruction using GCV and SURE-based methods," *IEEE Transactions on Image Processing*, vol. 21, no. 8, pp. 3659–3672, 2012.
- [22] F. Luisier, T. Blu, and M. Unser, "A new SURE approach to image denoising: Interscale orthonormal wavelet thresholding," *IEEE Transactions on Image Processing*, vol. 16, no. 3, pp. 593–606, 2007.
- [23] T. Blu and F. Luisier, "The SURE-LET approach to image denoising," *IEEE Transactions on Image Processing*, vol. 16, no. 11, pp. 2778–2786, 2007.
- [24] V. D. Ville and M. Kocher, "Nonlocal means with dimensionality reduction and SURE-based parameter selection," *IEEE Transactions on Image Processing*, vol. 20, no. 9, pp. 2683–2690, 2011.
- [25] K. Khare, C. J. Hardy, K. F. King, P. A. Turski, and L. Marinelli, "Accelerated MR imaging using compressive sensing with no free parameters," *Magnetic Resonance in Medicine*, vol. 68, no. 5, pp. 1450–1457, 2012.
- [26] C. Deledalle, S. Vaiter, J. M. Fadili, and Peyre, "Stein unbiased GrAdient estimator of the risk (SUGAR) for multiple parameter selection," *SIAM Journal on Imaging Sciences*, vol. 7, no. 4, pp. 2448–2487, 2014.
- [27] C. Guo and M. Davies, "Near optimal compressed sensing without priors: Parametric SURE approximate message passing," *IEEE Transactions on Signal Processing*, vol. 63, no. 8, pp. 2130–2141, 2015.
- [28] M. Lustig and J. M. Pauly, "SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space," *Magnetic Resonance in Medicine*, vol. 64, no. 2, pp. 457–471, 2010.
- [29] D. Liang, B. Liu, J. Wang, and L. Ying, "Accelerating SENSE using compressed sensing," *Magnetic Resonance in Medicine*, vol. 62, no. 6, pp. 1574–1584, 2009.
- [30] E. Ilicak, L. Senel, E. Biyik, and T. Çukur, "Profile-encoding reconstruction for multiple-acquisition balanced steady-state free precession imaging," *Magnetic Resonance in Medicine*, vol. 78, no. 4, pp. 1316–1329, 2017.
- [31] B. Bilgic, V. K. Goyal, and E. Adalsteinsson, "Multi-contrast reconstruction with bayesian compressed sensing," *Magnetic Resonance in Medicine*, vol. 66, no. 6, pp. 1601–1615, 2011.
- [32] J. Huang, C. Chen, and L. Axel, "Fast multi-contrast MRI reconstruction," *Magnetic Resonance Imaging*, vol. 32, no. 10, pp. 1344–1352, 2014.
- [33] E. Biyik, E. Ilicak, and T. Çukur, "Reconstruction by calibration over tensors for multi-coil multi-acquisition balanced SSFP imaging," *Magnetic Resonance in Medicine*, vol. 78, no. 4, pp. 1316–1329, 2017.
- [34] T. Hilbert, D. Nguyen, J.-P. Thiran, G. Krueger, T. Kober, and O. Bieri, "True constructive interference in the steady state (trueCISS)," *Magnetic Resonance in Medicine*, vol. 74, no. 4, pp. 1901–1910, 2018.
- [35] E. Gong, F. Huang, K. Ying, W. Wu, S. Wang, and C. Yuan, "Promise: Parallel-imaging and compressed-sensing reconstruction of multicontrast imaging using sharable information," *Magnetic Resonance in Medicine*, vol. 73, no. 2, pp. 523–535, 2015.
- [36] T. Kilic, E. Ilicak, T. Çukur, and E. U. Saritas, "Improved SPIRiT operator for joint reconstruction of multiple T2-weighted images," *Proceedings of the 25th Annual Meeting of ISMRM, Abstract no. 5165*, 2017.
- [37] B. Bilgic, T. H. Kim, C. Liao, M. Manhard, L. Wald, J. Haldar, and K. Setsompop, "Improving parallel imaging by jointly reconstructing multi-contrast data," *Magnetic Resonance in Medicine*, vol. 80, no. 2, pp. 619–632, 2018.
- [38] J. V. Velikina, A. L. Alexander, and A. Samsonov, "Accelerating MR parameter mapping using sparsity-promoting regularization in parametric dimension," *Magnetic Resonance in Medicine*, vol. 70, no. 5, pp. 1263–1273, 2013.
- [39] B. Zhao, W. Lu, T. K. Hitchens, F. Lam, C. Ho, and Z.-P. Liang, "Accelerated MR parameter mapping with low-rank and sparsity constraints," *Magnetic Resonance in Medicine*, vol. 74, no. 2, pp. 489–498, 2015.
- [40] T. Zhang, J. M. Pauly, and I. R. Levesque, "Accelerating parameter mapping with a locally low rank constraint," *Magnetic Resonance in Medicine*, vol. 73, no. 2, pp. 655–661, 2015.
- [41] M. Tofighi, K. Kose, and A. E. Cetin, "Denoising using projections onto the epigraph set of convex cost functions," *IEEE International Conference on Image Processing (ICIP)*, pp. 2709–2713, 2014.
- [42] A. Khare, U. Tiwary, W. Pedrycz, and M. Jeon, "Multilevel adaptive thresholding and shrinkage technique for denoising using Daubechies complex wavelet transform," *The Imaging Science Journal*, vol. 58, no. 6, pp. 340–358, 2010.
- [43] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projections onto the l_1 -ball for learning in high dimensions," *Proceedings of the 25th International Conference on Machine Learning*, pp. 272–279, 2008.
- [44] N. Parikh, "Proximal Algorithms," *Foundations and Trends in Optimization*, vol. 1, no. 3, pp. 127–239, 2014.
- [45] M. Tofighi, K. Kose, and A. E. Cetin, "Denoising images corrupted by impulsive noise using projections onto the epigraph set of the total variation function (PES-TV)," *Signal, Image and Video Processing*, vol. 9, no. 1, pp. 41–48, 2015.
- [46] E. Ilicak and T. Çukur, "Parameter-Free profile encoding reconstruction for multiple-acquisition bSSFP imaging," *Proceedings of the 25th Annual Meeting of ISMRM, Abstract no. 0571*, 2017.

- [47] M. Uecker, P. Lai, M. J. Murphy, P. Virtue, M. Elad, J. M. Pauly, S. S. Vasanawala, and M. Lustig, "ESPIRiT - an eigenvalue approach to autocalibrating parallel MRI: where SENSE meets GRAPPA," *Magnetic Resonance in Medicine*, vol. 71, no. 3, pp. 990–1001, 2014.
- [48] G. Chierchia, N. Pustelnik, J. Pesquet, and B. Pesquet-Popescu, "Epigraphical projection and proximal tools for solving constrained convex optimization problems," *Signal, Image and Video Processing*, vol. 9, no. 8, pp. 1737–1749, 2015.
- [49] J. Trzasko and A. Manduca, "Calibrationless parallel MRI using CLEAR," *Conference Record of the Forty Fifth Asilomar Conference on Signals, Systems and Computers (ASILOMAR)*, pp. 75–79, 2011.
- [50] P. J. Shin, P. E. Larson, M. A. Ohliger, M. Elad, J. M. Pauly, D. B. Vigneron, and M. Lustig, "Calibrationless parallel imaging reconstruction based on structured low-rank matrix completion," *Magnetic Resonance in Medicine*, vol. 72, no. 4, pp. 959–970, 2014.
- [51] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [52] K. Kose, V. Cevher, and A. Cetin, "Filtered variation method for denoising and sparse signal processing," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3329–3332, 2012.
- [53] F. Knoll, K. Bredies, T. Pock, and R. Stollberger, "Second order total generalized variation (TGV) for MRI," *Magnetic Resonance in Medicine*, vol. 65, no. 2, pp. 480–491, 2011.