

Accelerated Water–Fat Imaging Using Restricted Subspace Field Map Estimation and Compressed Sensing

Samir D. Sharma,* Houchun H. Hu, and Krishna S. Nayak

Water–fat separation techniques play an important role in a variety of clinical and research applications. In particular, multiecho separation methods remain a topic of great interest due to their ability to resolve water and fat images in the presence of B_0 -field inhomogeneity. However, these methods are inherently slow as they require multiple measurements. An accelerated technique with reduced k -space sampling is desirable to decrease the scan time. This work presents a new method for water–fat separation from accelerated multiecho acquisitions. The proposed approach does not require the region-growing or region-merging schemes that are typically used for field map estimation. Instead, the water, fat, and field map signals are estimated directly from the undersampled k -space measurements. In this work, up to 2.5 \times -acceleration is demonstrated in a water–fat phantom, ankle, knee, and liver. *Magn Reson Med* 67:650–659, 2012. © 2011 Wiley Periodicals, Inc.

Key words: water–fat separation; chemical-shift imaging; accelerated imaging; compressed sensing

Water–fat separation techniques play an important role in a variety of clinical and research applications. These methods have been used to detect fatty infiltration in the myocardium (1), to measure hepatic fat fraction (2), to image cartilage in the knee (3), and to identify brown adipose tissue (4). In particular, multiecho separation methods based on the work of Dixon (5) remain a topic of great interest due to their ability to resolve water and fat images in the presence of B_0 -field inhomogeneity. In these methods, k -space data are collected at multiple echo times (TEs) to resolve the water, fat, and field map signals. Separation occurs via a number of routes, one of which is the iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL) (6,7). This technique has become very popular due to its general formulation that allows for multiple chemical species and arbitrary TEs, while achieving the Cramer–Rao bound for an unbiased estimator (8). However, multiecho methods are inherently slow as measurements at multiple time points are required. This requirement limits the spatial resolution and volumetric coverage while increases the possibility of motion arti-

facts. An accelerated technique with reduced k -space sampling is desirable to decrease the total scan time of these multiecho acquisitions.

At the heart of multiecho methods lies estimation of the B_0 -field map inhomogeneity. The estimation of the field map is a challenging task as the least-squares cost function is nonconvex and many global minima may potentially exist. This ambiguity is typically resolved by incorporating the additional knowledge that the true field map is spatially smoothly varying. Many methods focusing on the estimate of the field map have been proposed, each essentially differing in how the presumed knowledge of smoothness is used. Yu et al. (9) developed a region-growing (RG) method in which the initial field map guess for the current pixel in question was determined by the field map values of previously estimated neighboring pixels. Berglund et al. have also recently developed a multiseed RG scheme (10). In their work, Lu and Hargreaves (11) first determine all the field map values for each pixel, within an expected range of field inhomogeneity, which would yield a local minimum in the least-squares cost function. They then use a field map growth scheme to assign estimates to each pixel while ensuring field map smoothness. Jacob and Sutton (12) take a conceptually similar route, although the details quickly diverge. Here, the possible field map values for each pixel are determined via a modification of the harmonic retrieval framework. The ambiguity is resolved by solving a cost function that promotes smoothness of the field map. A different approach has been proposed by Tsao and Jiang (13). Rather than considering each individual pixel, this work first treats each of echo-time images as one large pixel and estimates one field map value for all pixels. Subsequently, the images are decomposed into smaller regions and the field map value is updated for each region. This approach can be continued until the region size becomes one pixel. Last, recent work from Hernando et al. (14) aims to estimate the field map image on a finely discretized grid. Using the VARPRO formulation (15) for dimensionality reduction and a regularization term that penalizes nonsmooth field maps, a graph cut-based algorithm is used to find the desired field map. All the aforementioned works rely on knowledge of the fully k -space sampled multiecho images, which are not inherently obtained in an accelerated acquisition.

Ma et al. has applied SENSE parallel imaging (16) to accelerate multiecho acquisitions (17). Although straightforward, it is not clear how phase errors in the SENSE reconstruction would propagate to the separation step as preservation of phase information is critical. Reeder et al. (18) proposed to use the homodyne reconstruction

Ming Hsieh Department of Electrical Engineering, University of Southern California, Los Angeles, California, USA.

Grant sponsor: NIH; Grant numbers: R21DK081173, NIH K25DK087931.

*Correspondence to: Samir D. Sharma, M.S., University of Southern California, 3740 McClintock Ave., EEB400, Los Angeles, CA 90089.

E-mail: sdsharma@usc.edu

Received 9 December 2010; revised 5 May 2011; accepted 21 May 2011.

DOI 10.1002/mrm.23052

Published online 28 June 2011 in Wiley Online Library (wileyonlinelibrary.com).

technique for partial k -space acquisitions. The need for a fully sampled central k -space region necessarily limits this technique to acceleration rates of less than two. Recently, compressed sensing (CS) (19–21) has been used to accelerate water–fat imaging (22), where the water, fat, and field map images are solved simultaneously using sparsity constraints on each of the images. The field map is initialized by analytically finding the possible field map solutions from a fully sampled central k -space region, followed by a region-merging scheme to impose smoothness.

In the present work, a new method is proposed for water–fat separation from accelerated multiecho acquisitions. The traditional single-pixel model is generalized to consider estimation of the water, fat, and field map images directly from the undersampled k -space data. Similar to IDEAL, the estimation of water and fat images and the field map image is performed iteratively. Water and fat images are estimated via CS by exploiting their compressibility in the wavelet domain. In contrast to the work of Doneva et al. (22), the proposed approach does not require an initial seed voxel or a region-growing/region-merging scheme to impose field map smoothness. Instead, the field map estimate is updated in a restricted linear subspace. This restriction is enforced to avoid converging to local minima field map estimates, which may cause a water–fat swap. As we explain the field map update procedure in later sections, it may be helpful to remember that the l_1 -regularization term commonly found in CS reconstruction is not used to estimate the field map update term. The two steps (i.e. water–fat estimation and field map update) are iterated until the field map estimate converges. At this time, the linear subspace becomes less restrictive and the iterations are repeated. This process continues until a predefined stopping point is reached. We demonstrate up to 2.5 \times -acceleration in a water–fat phantom, ankle, knee, and liver.

THEORY

We begin with a mathematical description of the model, which is a generalization of the conventional single-pixel model. After establishing this, we describe in detail the water and fat image estimation as well as the approach for updating the field map estimate. We then conclude this section by summarizing the approach.

Signal Model

At echo-time t_n the single-pixel image-domain signal may be modeled as

$$s^p(t_n) = (\rho_w^p + (\rho_f^p \cdot d_n)) \cdot e^{j2\pi\psi^p t_n} \quad [1]$$

where ρ_w and ρ_f are, respectively, the unknown complex water and fat signals and ψ represents the unknown real-valued B_0 -field inhomogeneity (Hz). The superscript $p = 1:P$ denotes the pixel, the subscript n serves as an index over the TE points with $n = 1:N$, and the parameter d_n is a known complex-valued term resulting from the chemical shift between water and fat (23).

Extending the single-pixel model to an image model, we have

$$\mathbf{s}(t_n) = \mathbf{E}_n(\boldsymbol{\rho}_w + (\boldsymbol{\rho}_f \cdot d_n)) \quad [2]$$

where a bold typeface is used for matrix–vector notation and \mathbf{E}_n is a $P \times P$ diagonal matrix with $e^{j2\pi\psi^p t_n}$ on the p th diagonal.

Last, as the measurement data are undersampled k -space, we relate the unknowns to the measurements as

$$\mathbf{k}(t_n) = \Phi_u(\mathbf{E}_n(\boldsymbol{\rho}_w + (\boldsymbol{\rho}_f \cdot d_n))) \quad [3]$$

where Φ_u represents the undersampled Fourier transform (21) and $\mathbf{k}(t_n)$ represents the undersampled k -space at echo-time t_n . Equation 3 presents the mathematical description of the model that relates the unknowns to the measurement data. We do not model chemical-shift off-resonance (24) because the effects of this artifact are minimal in Cartesian acquisitions with a high sampling bandwidth.

Estimation

The unknown parameters are estimated via an iterative two-step process, similar to the method proposed by Reeder et al. (6). In the first step, the water and fat images are estimated while the current field map image estimate is maintained. We exploit the presumed compressibility of both the water and fat images to guide the reconstruction to the desired solution. In the next step, the estimate of the field map is updated. Rather than updating the field map estimate in the full-dimensional space, we restrict the update to a subspace to avoid falling into local minima solutions. The restriction on the subspace is gradually eased in further field map update steps.

Water–Fat Image Estimation

Fixing the field map image estimate ψ , Eq. 3 is linear in the unknown water and fat images. Additionally, MR images have been shown to be highly compressible in fixed transforms (25). Guided by the theory of CS, we estimate the water and fat images as

$$\min_{\boldsymbol{\rho}_w, \boldsymbol{\rho}_f} \left[\sum_{n=1}^N \|\mathbf{k}(t_n) - \Phi_u(\mathbf{E}_n(\boldsymbol{\rho}_w + \boldsymbol{\rho}_f \cdot d_n))\|_2^2 + \lambda \|\mathbf{W} \cdot \boldsymbol{\rho}_w\|_1 + \lambda \|\mathbf{W} \cdot \boldsymbol{\rho}_f\|_1 \right] \quad [4]$$

where \mathbf{W} represents an invertible transform in which the water and fat images are expected to be highly compressible and the parameter λ represents the tradeoff between the importance of data fidelity versus the prior information of image compressibility. This parameter can be different for $\mathbf{W} \cdot \boldsymbol{\rho}_w$ and $\mathbf{W} \cdot \boldsymbol{\rho}_f$, but we use the same one for simplicity while acknowledging the potential for further optimization.

By formulating and thus solving for images rather than individual pixels, we are able to incorporate knowledge of the images that we hope to reconstruct. The reconstruction problem posed in Eq. 4 is a convex function that can be reliably solved in a variety of ways (26).

Field Map Image Update

Estimating the field map is not as straightforward as estimating the water and fat images because the least-

squares cost function is nonconvex with respect to the field map image. Neighborhood information is often incorporated into the reconstruction to avoid erroneous field map estimates. Motivated by the work of Tsao and Jiang (13), we propose to restrict the dimension of the linear subspace in which the field map estimate is updated. The subspace dimension is gradually increased to allow for a more accurate field map estimate, but the dimension of the space always remains less than the number of voxels of the field map image. Further, the vectors that span the subspace are created in a way that promotes a spatially smoothly varying field map.

We first explain the approach conceptually with a one-dimensional (1D) example for a field map of 256 pixels. At first, strict restrictions are placed on the subspace in which the field map update is estimated. Figure 1a shows the one linear vector that spans the first instance of the restricted subspace. The single constant-valued vector restricts the field map to assuming only one value that is equal for all pixels. The field map is updated in this subspace until a convergence point is reached. Subsequently, the subspace is enlarged. Figure 1b depicts the three vectors that span this new space. The field map update at this step will thus be a superposition of these three vectors. Again, the field map update is estimated in this subspace until a convergence point is reached. The same procedure is repeated using the vectors spanned in Fig. 1c. Further steps continue the process of updating the field map in subspaces of increasing dimension until a predefined stopping point is reached. We defer details regarding the specific implementation (e.g. step size and stopping point) to the section Materials and Methods.

The 1D example extends in a quite straightforward manner to higher dimensions. In 2D, for instance, a vector is created via an outer-product of one vector from each dimension. Referring to the previous examples, the first subspace in a 2D problem would again restrict the field map estimate to the same value for each pixel, whereas the next step would result in a 2D subspace spanned by 9 ($=3^2$) 2D vectors.

Having presented the conceptual approach, we now turn to the mathematical details. Equation 3 is rewritten as

$$\mathbf{k}(t_n) = \Phi_u(\Delta \mathbf{E}_n \hat{\mathbf{E}}_n ((\hat{\rho}_w + \Delta \rho_w) + (\hat{\rho}_f + \Delta \rho_f) \cdot d_n)) \quad [5]$$

where, for example, $\hat{\rho}_w$ is the current (known) water image estimate and $\Delta \rho_w$ represents the error (unknown) in the estimate. The terms $\hat{\mathbf{E}}_n$ and $\Delta \mathbf{E}_n$ are defined in a similar manner as \mathbf{E}_n , but now with $\hat{\psi}^p$ and $\Delta \psi^p$ respectively. Using the first-order Taylor approximation for $\Delta \psi$ results in

$$\mathbf{k}(t_n) = \Phi_u(\Delta \mathbf{F}_n \hat{\mathbf{E}}_n ((\hat{\rho}_w + \Delta \rho_w) + (\hat{\rho}_f + \Delta \rho_f) \cdot d_n)) \quad [6]$$

where $\Delta \mathbf{F}_n$ is a $P \times P$ diagonal matrix with $1 + j2\pi \Delta \psi^p t_n$ on the p th diagonal. The first-order Taylor approximation is reasonably accurate for $|\Delta \psi \cdot t_n| = 1/8$ or field map updates of less than 20 Hz. This was generally found to be the case for the datasets in this work.




FIG. 1. 1D example to illustrate the vectors that span the restricted subspace in which the field map estimate is updated; one can imagine taking an inner product of the field map with each vector to obtain a coefficient in the restricted subspace; further, a backprojection would map the coefficient to image space. **a**: One vector that spans the subspace for the first instance of the field map update. **b**: Three vectors that span the subspace for the next instance of the field map update. **c**: Seven vectors that span the subspace in the third instance of the field map update; in this example, the step size factor is 0.5 as the vector in (a) has a support size of 256, the vectors in (b) each have a support size of 128, and the vectors in (c) each have a support size of 64; extensions to higher dimensions are made via outer-products of the 1D vectors.

Rearranging the known terms from the unknown terms and discarding $\Delta \cdot \Delta$ terms in a fashion similar to Reeder et al. (6), we have

$$\begin{aligned} \mathbf{k}(t_n) - \Phi_u(\hat{\mathbf{E}}_n(\hat{\rho}_w + \hat{\rho}_f \cdot d_n)) \\ = \Phi_u(\hat{\mathbf{E}}_n(\Delta \rho_w + \Delta \rho_f \cdot d_n) + (\Delta \mathbf{F}_n - \mathbf{I})\hat{\mathbf{E}}_n(\hat{\rho}_w + \hat{\rho}_f \cdot d_n)) \quad [7] \end{aligned}$$

At this point, we rewrite the field map update term, $\Delta \psi$, such that it is in the restricted subspace. We define \mathbf{R} as the matrix that contains the restricted subspace vectors. Each row of the matrix is one vector from that subspace. As an example and with reference to Fig. 1a, the first instance of \mathbf{R} would be a 1×256 matrix with each entry having the same value. The second instance of \mathbf{R} , with reference to Fig. 1b, would be a 3×256 matrix. The first row of \mathbf{R} would contain the values of the vector drawn with the solid line in Fig. 1b. Note that only the first 128 values of this vector are non-zero. The second and third rows of \mathbf{R} would contain, respectively, the values of the second and third vectors in Fig. 1b. In each new instance, \mathbf{R} is updated to reflect the current set of subspace vectors. For a given \mathbf{R} , the subspace coefficients are

$$\Delta \mathbf{m} = \mathbf{R} \cdot \Delta \psi \quad [8]$$

Using Eq. 8 to rewrite Eq. 7

$$\begin{aligned} \mathbf{k}(t_n) - \Phi_u(\hat{\mathbf{E}}_n(\hat{\rho}_w + \hat{\rho}_f \cdot d_n)) \\ = \Phi_u(\hat{\mathbf{E}}_n(\Delta \rho_w + \Delta \rho_f \cdot d_n) + \Delta \mathbf{M}_n \hat{\mathbf{E}}_n(\hat{\rho}_w + \hat{\rho}_f \cdot d_n)) \quad [9] \end{aligned}$$

where \mathbf{R}^H is the Hermitian transpose of \mathbf{R} and $\Delta \mathbf{M}_n$ is a $P \times P$ diagonal matrix containing the p th element of $j2\pi \cdot (\mathbf{R}^H \cdot \Delta \mathbf{m}) \cdot t_n$ on the p th diagonal. The unknowns are now $\Delta \rho_w$, $\Delta \rho_f$, and $\Delta \mathbf{m}$. By denoting the left side and right side of Eq. 9 as $\mathbf{r}(t_n)$ and $\mathbf{x}(\Delta \rho_w, \Delta \rho_f, \Delta \mathbf{m}; t_n)$, respectively, we arrive at the expression to estimate the field map update as well as estimates of the water and fat error images:

$$\min_{\Delta\rho_w, \Delta\rho_f, \Delta\mathbf{m}} \left[\sum_{n=1}^N \|\mathbf{r}(t_n) - \mathbf{x}(\Delta\rho_w, \Delta\rho_f, \Delta\mathbf{m}; t_n)\|_2^2 + \lambda |\mathbf{W} \cdot \Delta\rho_w|_1 + \lambda |\mathbf{W} \cdot \Delta\rho_f|_1 \right] \quad [10]$$

As in IDEAL, the field map update $\Delta\psi = \mathbf{R}^H \cdot \Delta\mathbf{m}$ is added to the current field map estimate while the other estimates are discarded. Again, for simplicity, the λ and \mathbf{W} used in Eq. 10 is the same as used in Eq. 4.

Summary of Estimation Procedure

Figure 2 illustrates the proposed procedure, which is summarized below.

1. Initialize field map estimate to 0 and \mathbf{R} to contain one constant-valued vector.
2. Estimate water and fat images (Eq. 4).
3. Update the field map estimate within \mathbf{R} (Eq. 10).
4. IF mean of the absolute field map update is small (e.g. < 1 Hz)
 - Go to Step 5;
 - ELSE
 - Go to Step 2;
5. IF predefined maximum subspace size is reached
 - Done;
 - ELSE
 - Update \mathbf{R} to span a larger subspace;
 - Go to Step 2.

MATERIALS AND METHODS

Data Collection

Experiments were conducted on a 3-T Signa EXCITE HDx system (GE Healthcare, Waukesha, WI). Fully sampled data were collected with volunteer consent using a GE-investigational IDEAL 3D spoiled-gradient-echo sequence at three TE points ($TE_1 = 2.184$ ms, $\Delta TE = 0.794$ ms) with one echo per TR, BW = 125 kHz, a 256×256 sampling matrix, and a flip angle of 5° .

Data were first acquired from a water-fat phantom using a single-channel birdcage coil. An oblique coronal slab through the water-fat interface captured the spectrum from entirely water to entirely fat. The field-of-view (FOV) was 15 cm and slice thickness (Δz) was 6 mm. Data were also acquired from the ankle (FOV = 24 cm and $\Delta z = 5$ mm) and from the knee (FOV = 18 cm and $\Delta z = 5$ mm) with the single-channel coil. Using an eight-channel knee transmit/receiver array, data were collected from the knee (FOV = 18 cm and $\Delta z = 5$ mm) and an eight-channel torso receiver array was used to collect a liver dataset (FOV = 34 cm and $\Delta z = 5$ mm).

Estimation

The fully sampled data were retrospectively down-sampled to simulate an accelerated acquisition. Down-sampling was done in a 256×3 (number of phase encodes \times number of TEs) space. This accelerated acquisition scheme could be achieved prospectively by blipping the phase encoding gradient between subsequent TEs so that a different phase encoding line is

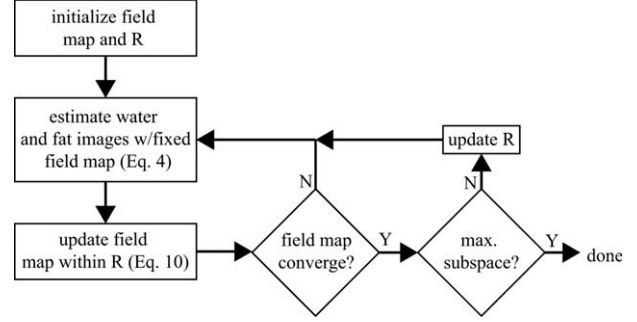


FIG. 2. Flowchart of the proposed estimation process; the field map is initialized to zero and the restricted subspace \mathbf{R} is initialized as in Fig. 1a; the water and fat images are estimated with the fixed field map; the field map is then updated in the restricted subspace \mathbf{R} ; this process is continued until the field map update is small; \mathbf{R} is then eased to span a larger dimension and the entire process is repeated until a predefined stopping point is reached.

acquired at each TE point. This would result in an undersampled k_y - TE space. We used a variable-density sampling scheme that minimized the maximum sidelobe of the point-spread function in an effort to suppress aliasing interference (21). The central 1/8th k -space phase-encode lines were sampled at all TEs. The Daubechies-8 wavelet transform was used as the sparsifying basis for the water and fat images. Previous experiments had shown that this transform provided an accurate approximation of water and fat images from only 15–25% of the highest magnitude coefficients. We used the six-peak fat model proposed by Yu et al. (23). Last, in Eqs. 4 and 10, the parameter λ was set to 2 for $2\times$ -acceleration and 3 for $2.5\times$ -acceleration for every reconstruction presented in this work. These values were determined empirically by comparing the reconstruction quality of the water and fat images with fully sampled reference images.

The subspace dimension for the field map update was gradually increased such that the support size of each 1D vector at step m was $3/4$ th the size of the support in step $m - 1$. For example, at step 1, there was one constant-valued vector with support size of 256. At the next step, each vector had a support size of 192 ($=256 \times 0.75$), then 144 ($=192 \times 0.75$), and so on. Noninteger results were rounded to the nearest integer. The number of vectors that spanned a restricted subspace was calculated as follows. If the support size S of each subspace vector was a power of 2, then the number of vectors was

$$\frac{512}{S} - 1$$

For example, if $S = 256$, then the number of vectors is one, as is the case in Fig. 1a. If $S = 128$, then the number of vectors is three as in Fig. 1b. For support sizes that were not a power of 2, the number of vectors was calculated as the rounded average of the number of vectors from the two closest powers of 2. For example, for a support size of 144, the number of vectors in the corresponding subspace was two, which is the average of one

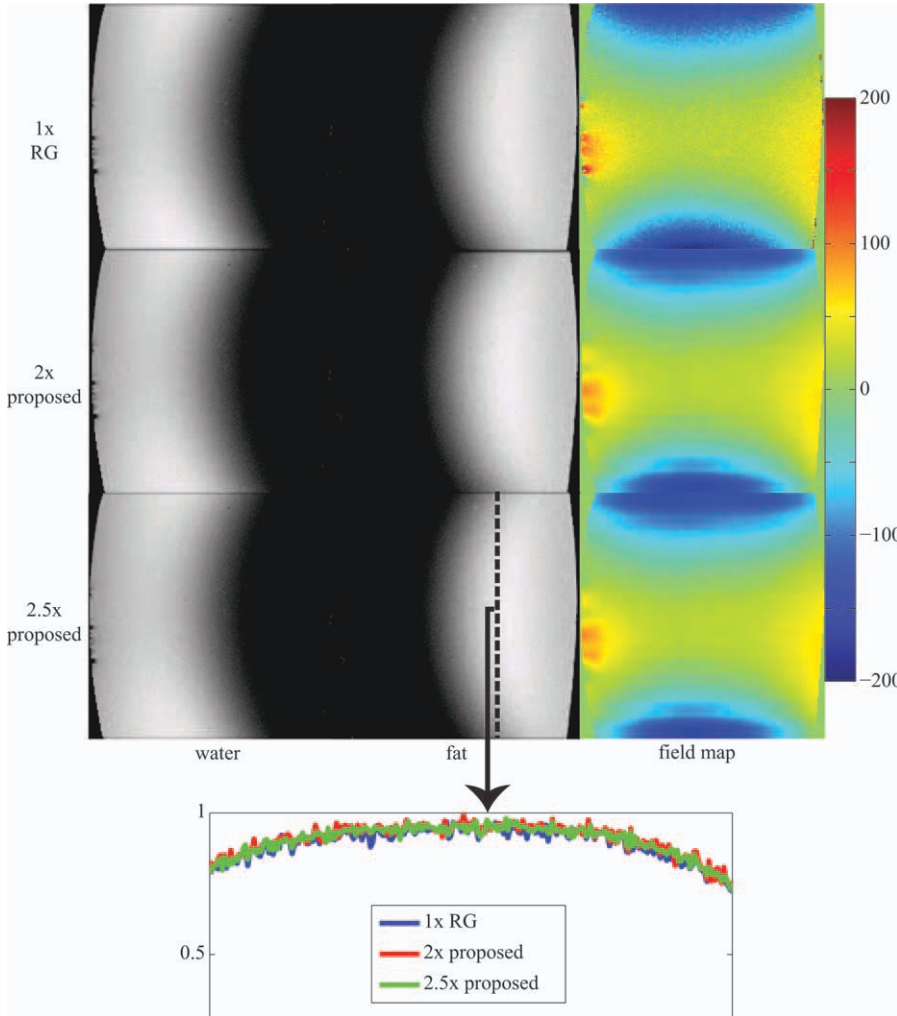


FIG. 3. Image estimates from a 256×256 water-fat phantom using a single-channel birdcage coil; estimates using the proposed method at $2\times$ and $2.5\times$ are visually similar to those from $1\times$ RG; the plot at the bottom shows a profile through the fat image taken at the black dashed line; field map colorbar is in Hertz (Hz). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

and three that result from a $S = 256$ and $S = 128$, respectively. The number of vectors was calculated in this manner to ensure overlapping supports. As described earlier, these vectors were created in 1D and expanded to 2D via outer-products.

The iterative estimation procedure within one subspace continued until the mean absolute value of the current field map update was less than 1 Hz. Finally, the entire estimation routine continued until the 1D vector support size was less than 1/16th of the image dimension. For a 256×256 image, the stopping point occurred when the support size was less than 16 pixels.

For multicoil acquisitions, each coil was processed separately. The individual field maps were then combined by

$$\psi_c^p = \frac{\sum_{a=1}^A (|\hat{\rho}_w^a|^2 + |\hat{\rho}_f^a|^2)^{.a} \psi^p}{\sum_{b=1}^B (|\hat{\rho}_w^b|^2 + |\hat{\rho}_f^b|^2)} \quad [11]$$

where $A = B =$ the number of coils and a left superscript denotes the coil element. Water and fat images were then re-estimated for each coil using the combined field map. The results were combined using root-sum-of-squares.

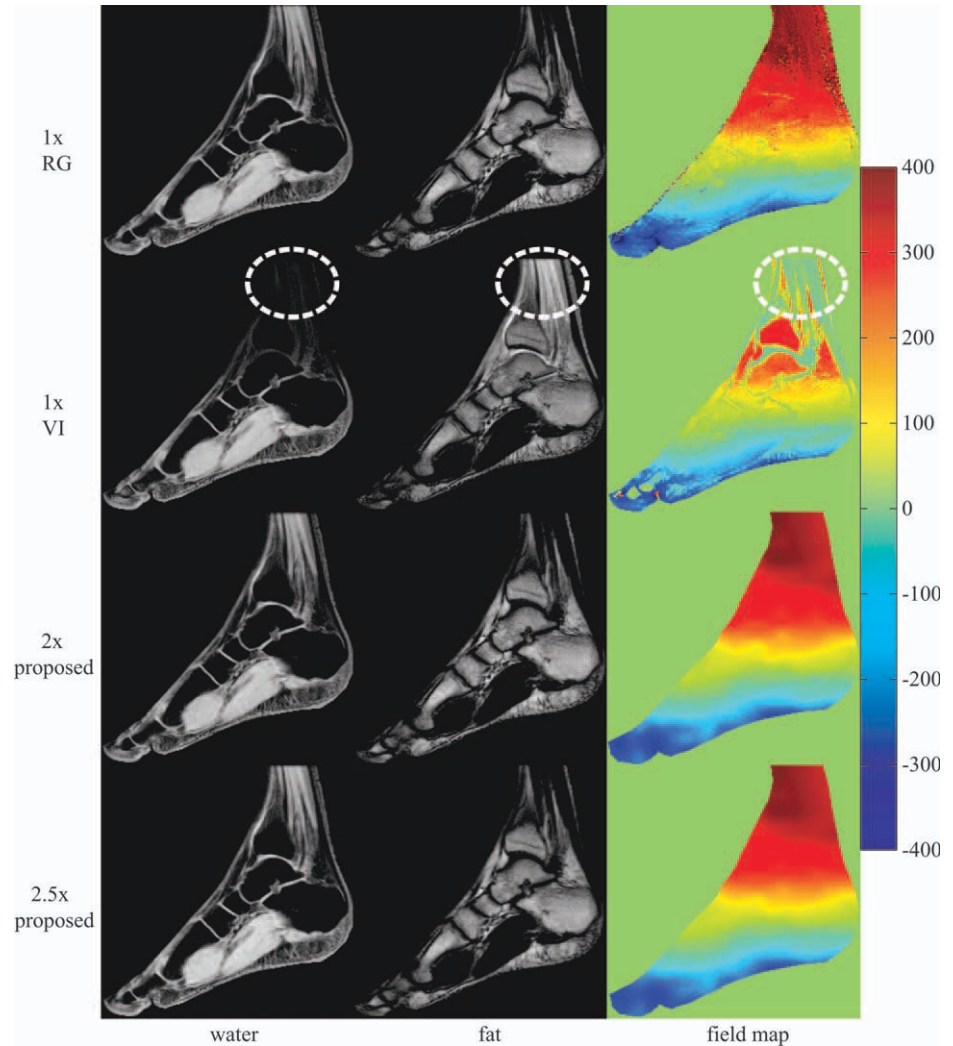
The approach was tested at $2\times$ and $2.5\times$ on a water-fat phantom, ankle, knee, and liver. Results were compared against a nonaccelerated voxel-independent IDEAL (IDEAL-VI) method (6) and/or an in-house implementation of IDEAL with RG (IDEAL-RG), based on the work of Yu et al. (9). All processing was done offline in Matlab (The Mathworks, Inc, Natick, MA) using a conjugate-gradient implementation (21) and Wavelab for the wavelet transform (27).

RESULTS

Figure 3 displays 256×256 water, fat, and field map images that were estimated using IDEAL-RG and the proposed method at $2\times$ - and $2.5\times$ -acceleration as well as a profile of signal intensity through a cross-section of the fat image. The proposed method accurately separates the water and fat signals at both acceleration rates. Although the field map image that is estimated using the proposed method appears smoother than that from IDEAL-RG, the estimated water and fat intensities are mostly unaffected as shown in the profile plot.

Figure 4 displays water, fat, and field map image estimates for a sagittal slice of the ankle using IDEAL-VI, IDEAL-RG, and the proposed method at $2\times$ - and $2.5\times$ -

FIG. 4. Image estimates from a 256×256 ankle using a single-channel birdcage coil; inaccurate field map estimates in the $1 \times$ VI method cause water-fat swaps (see outlined regions in VI images); estimates using the proposed method at $2 \times$ and $2.5 \times$ avoid local minima field map solutions resulting in proper water-fat separation; field map colorbar is in Hertz (Hz). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]



acceleration. The region outlined in the IDEAL-VI field map estimate illustrates an area where the field map was incorrectly estimated, causing inaccurate water-fat separation. Both IDEAL-RG and the proposed approach at $2 \times$ - and $2.5 \times$ -acceleration resolve the field map ambiguity, resulting in correct water-fat separation.

Figure 5 displays the estimates for a sagittal slice of the knee using IDEAL-RG and the proposed approach at $2 \times$ - and $2.5 \times$ -acceleration. Estimates from IDEAL-VI are not shown as they were nearly identical to those from IDEAL-RG. Figures 6 and 7, respectively, display the image estimates for a multicoil acquisition from an axial slice of the knee and from an axial slice of the liver. For the liver image, the data were acquired in one 38 s breath-hold using an eight-channel torso array. The proposed method at both acceleration rates accurately separates the water and fat images with no signal swaps. The water and fat image estimates from the proposed method at $2 \times$ -acceleration appear nearly identical to the corresponding fully sampled IDEAL-RG estimates, while there is only a slight hint of blurring in the image estimates from the proposed method at $2.5 \times$ -acceleration. The total computational time for reconstructing the water, fat, and

field map images was ~ 20 min (24 CPUs, 2.93 GHz, and 48 GB RAM).

DISCUSSION

The field map estimation approach that we have proposed avoids the local minima solutions that plague VI methods. This is most evident in the ankle dataset. Figure 8a shows the evolution of the field map estimate for a voxel that was incorrectly estimated using IDEAL-VI (the voxel is denoted by the black dot in Fig. 8b,c). With an initialization of 0 Hz, a descent-based VI method converges to -52.58 Hz, which is denoted by the red line in Fig. 8a. Using the proposed method, the field map is correctly estimated for that voxel. In fact, it appears that as early as the second instance of the restricted subspace (i.e. 1D support size = 192), the field map estimate is quite close to its true value. This may suggest that larger steps between the 1D support size could be taken, say with a step size factor of 0.5; however, we have observed much slower convergence in other datasets, which necessitated the current step size factor of 0.75. Figure 8b,c shows the field map image estimate at two different

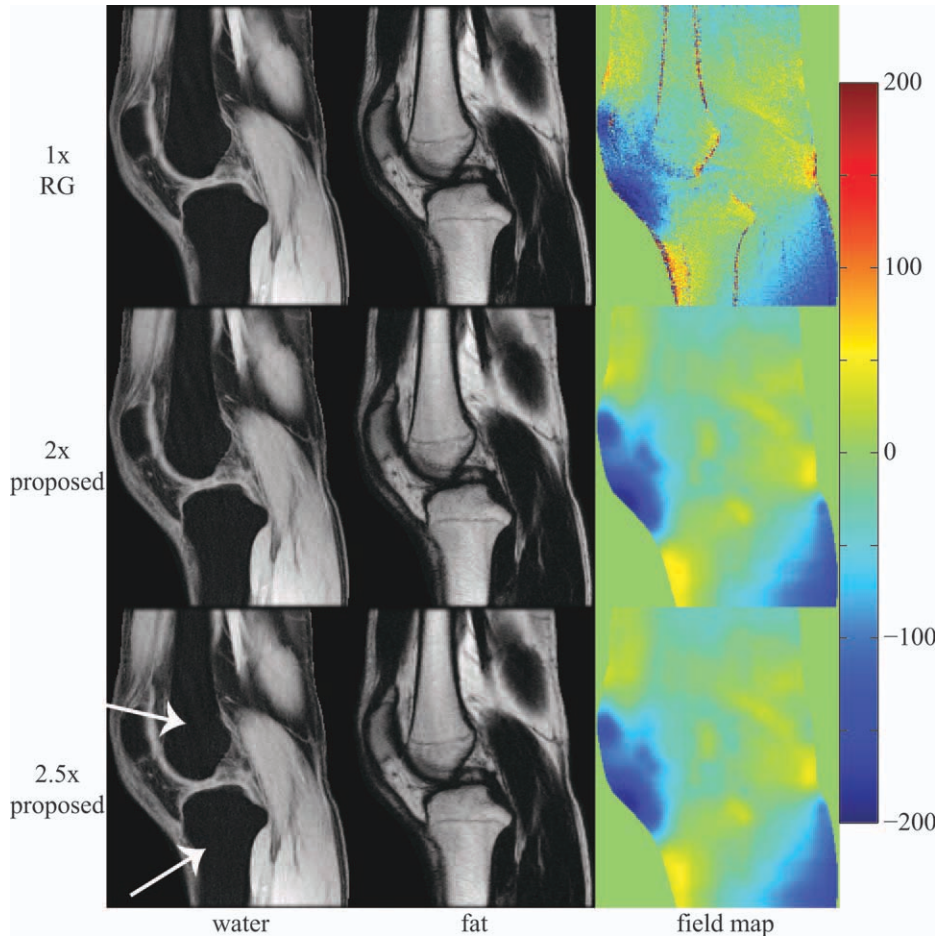


FIG. 5. Image estimates from a 256×256 knee using a single-channel birdcage coil; estimates at both $2\times$ and $2.5\times$ from the proposed method are very similar to those from $1\times$ RG; the white arrows in the figure highlight areas of slight discrepancy between estimates from RG and the proposed method; field map colorbar is in Hertz (Hz). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

points in the reconstruction process. Even in Fig. 8b, the field map estimate is quite close to the true value.

Generally, the field map estimate from the proposed method appeared smoother than that from IDEAL-RG. This is probably a direct result of using restricted subspaces for the field map updates. By further increasing the dimension of the subspace, we expect the field map estimates to approach those seen from IDEAL-RG. That said, the quality of the water–fat separation does not seem to be adversely affected by the stopping criterion that we chose in this work. This is supported by the profile plot of the fat image from the water–fat phantom in Fig. 3.


Updating the field map in a restricted subspace significantly reduces the number of coefficients that must be estimated. This becomes extremely convenient when estimating the field map from undersampled measurements. For example, the estimation of the field map image from the ankle was finally computed with 529 ($=23^2$) vectors in \mathbf{R} . This corresponds to using 1D vectors with a support size of 20 pixels. With an image size of 256×256 , this represents a “compression-ratio” of 0.8%. In 3D imaging at 256^3 pixels, the number of coefficients to calculate would be 12,167 ($=23^3$), which is a ratio of 0.07%. By restricting the field map estimate to particular coefficients, the proposed method may achieve a greater acceleration than those approaches that impose l_1 -regularization on the field map (22). This is because l_1 -regulari-

zation methods must pay a penalty of additional measurements as the locations of significant coefficients are unknown before estimation. By contrast, the proposed method specifies which coefficients to estimate (i.e. those coefficients in the restricted subspace).

While the restricted subspace approach provides convenience in estimation, it limits the types of field maps that one should hope to recover. Although off-resonance due to B_0 -field inhomogeneity tends to be smoothly varying, susceptibility-induced variations may cause sharp changes in the local field map, which would violate our assumption of smoothness. These variations are especially pronounced at air–tissue interfaces, which may occur in abdominal imaging. Another limitation in this work is the lack of compensation for R_2^* effects. This becomes especially critical in quantitative applications such as measuring hepatic fat fraction where high iron content would significantly attenuate the signal due to its short T_2^* value. The proposed model allows for the inclusion of this factor but care must be taken in estimating its value because, as with the field map, the least-squares cost function is nonconvex with respect to this parameter. Incorporation of T_2^* into the proposed model is under investigation.

A few user-defined design choices were necessary in this method. As with most other applications of CS, the regularization parameter λ must be set. This parameter was empirically determined, but was kept constant for

FIG. 6. Image estimates from a 256×256 knee using an eight-channel knee coil; estimates at $2\times$ and $2.5\times$ using the proposed method are very similar to the $1\times$ RG results; field map colorbar is in Hertz (Hz). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]



all 256×256 image reconstructions at a given acceleration rate. How this parameter would change as a function of image size, anatomy, and acceleration rate is an open question in this work as well as any other that uses CS reconstruction. The step size factor for updating the field map subspace was also empirically determined. The 0.75 factor worked for all datasets presented in this

article although it was overly conservative for some such as the water-fat phantom and the knee.

The design of subspace vectors was motivated mostly by simplicity. Field maps are generally not piecewise-linear as would be reflected by these pyramidal vectors. However, as the support size of each 1D vector decreases, the resulting estimate more closely




FIG. 7. Image estimates from a 256×256 liver using an eight-channel torso array; estimates from the $2\times$ and $2.5\times$ proposed method are very similar to the $1\times$ RG estimates; field map colorbar is in Hertz (Hz). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

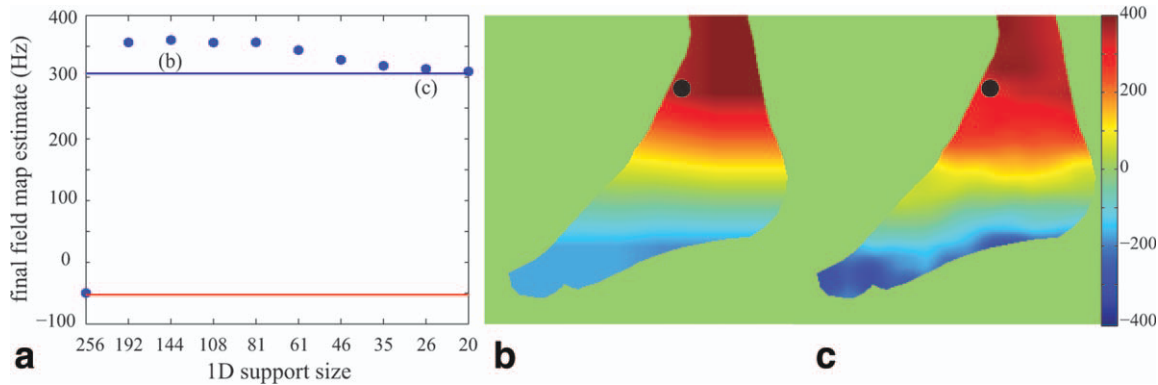


FIG. 8. **a**: The field map estimate at $2\times$ -acceleration for the voxel denoted by the black dot in **(b)** and **(c)** as a function of 1D support size; the red line and blue line indicate the field map estimate using $1\times$ IDEAL-VI and $1\times$ IDEAL-RG, respectively; **(b,c)** the evolution of the field map image estimate at two different support sizes as denoted in **(a)**; the two field map images use the same colorbar, which is in units of Hertz (Hz). [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

approximates the smooth variations of the field map. Subspace vectors that more accurately resemble the expected field map may be used to decrease the step size factor and provide a more accurate estimation. One possibility for learning these subspace vectors would be to first develop a set of characteristics for the sequence of spaces in which the field map will be updated. These characteristics should include elements as promoting a smoothly varying field map and preventing local minima field map solutions. The vectors that span each of these subspaces can then be learnt such that they provide a compact representation of the space.

Last, this work does not rely on any coil sensitivity information for a reduction of acquisition time. Parallel imaging techniques use complementary knowledge to that used in this work; thus, we do expect further acceleration by incorporating sensitivity information into the proposed accelerated water-fat imaging framework. This remains an item of future work.

CONCLUSIONS

In conclusion, we have presented a new approach for accelerated water-fat imaging. By modeling the images rather than single pixels, we were able to exploit presumed knowledge of the water, fat, and field map images to estimate their values from undersampled k -space data.

ACKNOWLEDGMENTS

The authors would like to thank Huanzhou Yu and Ann Shimakawa from GE Healthcare for technical support with the IDEAL sequence.

REFERENCES

- Kellman P, Hernando D, Shah S, Zuehlsdorff S, Jerecic R, Mancini C, Liang Z, Arai A. Multiecho Dixon fat and water separation method for detecting fibrofatty infiltration in the myocardium. *Magn Reson Med* 2009;61:215–221.
- Kim H, Taksali S, Dufour S, Befroy D, Goodman T, Petersen K, Shulman G, Caprio S, Constable R. Comparative MR study of hepatic fat quantification using single-voxel proton spectroscopy, two-point Dixon and three-point IDEAL. *Magn Reson Med* 2008;59:521–527.
- Kijowski R, Blankenbaker D, Woods M, Shinki K, De Smet A, Reeder S. 3.0-T evaluation of knee cartilage by using three-dimensional IDEAL GRASS imaging: comparison with fast spin-echo imaging. *Radiology* 2010;255:117–127.
- Hu H, Smith DJr, Nayak K, Goran M, Nagy T. Identification of brown adipose tissue in mice with fat-water IDEAL MRI. *J Magn Reson Imaging* 2010;31:1195–1202.
- Dixon W. Simple proton spectroscopic imaging. *Radiology* 1984;153:189–194.
- Reeder S, Wen Z, Yu H, Pineda A, Gold G, Markl M, Pelc N. Multi-coil Dixon chemical species separation with an iterative least-squares estimation method. *Magn Reson Med* 2004;51:35–45.
- Reeder S, Pineda A, Wen Z, Shimakawa A, Yu H, Brittain J, Gold G, Beaulieu C, Pelc N. Iterative decomposition of water and fat with echo asymmetry and least-squares estimation (IDEAL): application with fast spin-echo imaging. *Magn Reson Med* 2005;54:636–644.
- Pineda A, Reeder S, Wen Z, Pelc N. Cramer-Rao bounds for three-point decomposition of water and fat. *Magn Reson Med* 2005;54:625–635.
- Yu H, Reeder S, Shimakawa A, Brittain J, Pelc N. Field map estimation with a region growing scheme for iterative 3-point water-fat decomposition. *Magn Reson Med* 2005;54:1032–1039.
- Berglund J, Johansson L, Ahlström H, Kullberg J. Three point Dixon method enables whole body water and fat imaging of obese subjects. *Magn Reson Med* 2010;63:1659–1668.
- Lu W, Hargreaves B. Multiresolution field map estimation using golden section search for water-fat separation. *Magn Reson Med* 2008;60:236–244.
- Jacob M, Sutton B. Algebraic decomposition of fat and water in MRI. *IEEE Trans Med Imaging* 2009;28:173–184.
- Tsao J, Jiang Y. Hierarchical IDEAL: robust water-fat separation at high field by multiresolution field map estimation. In: *Proceedings of the 18th Annual Meeting of ISMRM*, Toronto, ON, Canada, 2008. p 653.
- Hernando D, Kellman P, Haldar J, Liang Z. Robust water/fat separation in the presence of large field inhomogeneities using a graph cut algorithm. *Magn Reson Med* 2010;63:79–90.
- Hernando D, Haldar J, Sutton B, Ma J, Kellman P, Liang Z. Joint estimation of water/fat images and field inhomogeneity map. *Magn Reson Med* 2008;59:571–580.
- Pruessmann K, Weiger M, Scheidegger M, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med* 1999;42:952–962.
- Ma J, Bankson J, Stafford R. Multipoint Dixon imaging using sensitivity encoding. In: *Proceedings of the 13th Annual Meeting of ISMRM*, Toronto, ON, Canada, 2003. p 1069.
- Reeder S, Hargreaves B, Yu H, Brittain J. Homodyne reconstruction and IDEAL water-fat decomposition. *Magn Reson Med* 2005;54:586–593.
- Candes E, Tao T. Near-optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans Inform Theory* 2006;52:5406–5425.
- Donoho D. Compressed sensing. *IEEE Trans Inform Theory* 2006;52:1289–1306.

21. Lustig M, Donoho D, Pauly J. Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn Reson Med* 2007;58:1182–1195.
22. Doneva M, Börnert P, Eggers H, Mertins A, Pauly J, Lustig M. Compressed sensing for chemical shift-based water-fat separation. *Magn Reson Med* 2010;64:1749–1759.
23. Yu H, Shimakawa A, McKenzie C, Brodsky E, Brittain J, Reeder S. Multiecho water fat separation and simultaneous $R2^*$ estimation with multifrequency fat spectrum modeling. *Magn Reson Med* 2008;60:1122–1134.
24. Brodsky E, Holmes J, Yu H, Reeder S. Generalized k -space decomposition with chemical shift correction for non-Cartesian water fat imaging. *Magn Reson Med* 2008;59:1151–1164.
25. Liang D, Liu B, Wang J, Ying L. Accelerating SENSE using compressed sensing. *Magn Reson Med* 2009;62:1574–1584.
26. Boyd S, Vandenberghe L. *Convex optimization*. Cambridge, United Kingdom: Cambridge University Press; 2004.
27. Buckheit J, Chen S, Donoho D, Johnstone I, Scargle J. *About Wavelab*. Wavelab software package version 0700, Technical Report. Stanford, CA: Department of Statistics, Stanford University; 1995.