

Low-dimensional-Structure Self-Learning and Thresholding: Regularization Beyond Compressed Sensing for MRI Reconstruction

Mehmet Akçakaya,^{1*} Tamer A. Basha,¹ Beth Goddu,¹ Lois A. Goepfert,¹
Kraig V. Kissinger,¹ Vahid Tarokh,³ Warren J. Manning,^{1,2} and Reza Nezafat¹

An improved image reconstruction method from undersampled k -space data, low-dimensional-structure self-learning and thresholding (LOST), which utilizes the structure from the underlying image is presented. A low-resolution image from the fully sampled k -space center is reconstructed to learn image patches of similar anatomical characteristics. These patches are arranged into “similarity clusters,” which are subsequently processed for dealiasing and artifact removal, using underlying low-dimensional properties. The efficacy of the proposed method in scan time reduction was assessed in a pilot coronary MRI study. Initially, in a retrospective study on 10 healthy adult subjects, we evaluated retrospective undersampling and reconstruction using LOST, wavelet-based l_1 -norm minimization, and total variation compressed sensing. Quantitative measures of vessel sharpness and mean square error, and qualitative image scores were used to compare reconstruction for rates of 2, 3, and 4. Subsequently, in a prospective study, coronary MRI data were acquired using these rates, and LOST-reconstructed images were compared with an accelerated data acquisition using uniform undersampling and sensitivity encoding reconstruction. Subjective image quality and sharpness data indicate that LOST outperforms the alternative techniques for all rates. The prospective LOST yields images with superior quality compared with sensitivity encoding or l_1 -minimization compressed sensing. The proposed LOST technique greatly improves image reconstruction for accelerated coronary MRI acquisitions. Magn Reson Med 66:756–767, 2011. ©2011 Wiley-Liss, Inc.

Key words: compressed sensing; accelerated imaging; image reconstruction; block matching; learning algorithm; cardiac MR; coronary MRI

Over the past two decades, several approaches have been proposed to improve image acquisition speed (1–9). Partial Fourier methods use Hermitian symmetry to reduce scan time (1). Non-Cartesian sampling techniques with more efficient k -space transversals have incoherent or less visually significant artifacts compared with Cartesian sampling

(2–6). The local sensitivity of phased array coil elements are used in parallel imaging either in k -space (simultaneous acquisition of spatial harmonics, SMASH (7) and generalized autocalibrating partially parallel acquisition, GRAPPA (9)) or image space (sensitivity encoding, SENSE (8)) to allow accelerated acquisition. Model-based techniques (10–12) and time-frequency approaches taking advantage of spatiotemporal correlations (4,13) have also been proposed to accelerate image acquisition in dynamic imaging. Despite all these efforts, long scan time is still a main challenge, especially in cardiac imaging where alternative modalities such as multidetector CT benefits from rapid acquisition and high patient throughput.

Compressed (or compressive) sensing (CS) is a recent image reconstruction approach for incoherent undersampling patterns (achieved by random sampling of k -space data in Cartesian acquisitions) that exploits the sparsity (or more generally compressibility) of the image in a transform domain (14,15) and may be used to surpass the current rapid acquisition techniques in terms of acceleration rate (15,16). CS reconstruction aims to maximize the sparsity of the reconstructed image in a transform domain subject to data consistency constraints comparing the estimate to the acquired k -space data. This is typically done by minimizing the convex l_1 norm of the transform domain coefficients (17,18). This accelerated technique has already been applied in several cardiac MRI applications to reduce scan time, which can be traded off for higher spatial or temporal resolution (19,20). Furthermore, attempts have been made to combine parallel imaging and CS, via sparsity-regularized iterative GRAPPA-type approaches (21,22), a concurrent combination of CS and SENSE where the coil sensitivities are used to enforce data consistency in addition to sparsity constraints (20), and a serial combination of CS and SENSE (23).

Even though image acceleration is possible using these methods, new reconstruction strategies that address the limitations of previous methods or enable higher acceleration rate are desirable. Partial Fourier and parallel imaging methods suffer from noise amplification, where the noise level and artifacts increase with the acceleration rate, to a level that might hinder clinical use. This is particularly important for higher spatial resolution applications, such as coronary MRI. Non-Cartesian techniques provide image acceleration with less coherent artifacts. However, improved reconstruction techniques, such as CS, which exploit the underlying image sparsity, can further reduce noise-like artifacts (5,14). Despite, the

¹Department of Medicine, Beth Israel Deaconess Medical Center and Harvard Medical School, Boston, Massachusetts, USA.

²Department of Radiology, Beth Israel Deaconess Medical Center and Harvard Medical School, Boston, Massachusetts, USA.

³School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA.

Grant sponsor: Harvard Clinical and Translational Science Center (National Center for Research Resources); Grant numbers: NIH R01EB008743-01A2, AHA SDG-0730339N, NIH UL1 RR025758-01.

*Correspondence to: Mehmet Akçakaya, Ph.D., Beth Israel Deaconess Medical Center, 330 Brookline Ave, Boston, MA 02215. E-mail: makcakay@bidmc.harvard.edu

Received 7 October 2010; revised 16 December 2010; accepted 3 January 2011.

DOI 10.1002/mrm.22841

Published online 4 April 2011 in Wiley Online Library (wileyonlinelibrary.com).

© 2011 Wiley-Liss, Inc.

improvement in noise-reduction, CS-based techniques still suffer from residual artifacts and image smoothing, because they assume an image has a sufficiently sparse representation in a preselected transform domain. Although sparsity is a necessary condition for l_1 norm reconstruction, it is not possible to know whether a transform can efficiently represent the underlying image characteristics. For instance, wavelets cannot capture smooth transitions sparsely, whereas finite differences have problems with sharp edges. Thus, minimization of sparsity with a fixed transform domain may result in blurring and other image artifacts. Even if the transform dictionary is generated using the computationally expensive dictionary learning algorithms (24), the effectiveness of such transforms may degrade due to inter- and intra-patient variability, especially for cardiac MR where contrast and signal level varies significantly between different acquisitions.

In this study, we sought to develop an improved CS reconstruction method from undersampled k -space data that utilizes the structure and anatomical features in the image being reconstructed. By using the information from the image itself, we aim to represent various features of the image sparsely in an adaptive fashion, without the need for training data. Once such an adaptive representation is achieved, aliasing artifacts and noise can be removed using thresholding approaches. Furthermore, learning the structure in this self-contained manner ensures robustness to variations among patients, exams, and contrast level. After the introduction of the reconstruction method, we present a comparative study between wavelet-based l_1 minimization, total variation (TV) regularization, and the proposed method for accelerated coronary MRI in a cohort of healthy subjects. We will also present results from prospectively acquired coronary MRI datasets, comparing the proposed method with the aforementioned techniques and with SENSE.

THEORY

Let $\mathbf{F}$ denote the unitary Fourier transform and $\mathbf{F}_\Omega$ be the operator that undersamples the k -space with the pattern Ω . Hence, the j th coil image, $\mathbf{m}_j(x, y, z)$ is measured as $\mathbf{S}_j = \mathbf{F}_\Omega(\mathbf{m}_j) + \mathbf{n}_j$ in the presence of additive noise. To estimate the image, CS reconstruction solves a minimization problem based on an objective function:

$$\hat{\mathbf{m}} = \arg \min_{\mathbf{m}} \frac{1}{2} \|\mathbf{S} - \mathbf{F}_\Omega(\mathbf{m})\|_2^2 + \tau \Phi(\Psi \mathbf{m}), \quad [1]$$

where the first term is a fidelity measure of image consistency, and the second term is a scalar weight τ of the sparsity regularizer Φ with a transform domain Ψ . Typically, Φ is chosen as the ℓ_p ($p \leq 1$) norm of the transform domain coefficients, which captures the sparsity of the image in transform domain Ψ . This equation is usually solved using an iterative thresholding algorithm, which alternates between enforcing k -space data consistency and dealiasing (25). At iteration t , image estimates of individual coils, $\mathbf{m}_j^{(t)}$ are Fourier transformed, and the estimated k -space lines are replaced with the acquired k -space lines $\mathbf{S}_j$. These new k -spaces are subsequently inverse Fourier

transformed to generate $\mathbf{v}_j^{(t)}$ in the data-consistency stage. In the dealiasing stage, the consistent images $\mathbf{v}_j^{(t)}$ are dealiased to generate $\mathbf{m}_j^{(t+1)}$. Conventionally, this is done by transforming $\mathbf{v}_j^{(t)}$ to Ψ domain, thresholding these coefficients individually and inverse Ψ -transforming to generate $\mathbf{m}_j^{(t+1)}$ (14,15,19). In this work, we replace this technique with an alternative dealiasing strategy, low-dimensional-structure self-learning and thresholding (LOST), which is described in detail below.

LOST learns the image areas of similar signal characteristics and uses this information for reconstruction. In the initial stage of the algorithm, 2D image patches of similar signal characteristics are grouped together into “similarity clusters.” These data are arranged into a 3D structure and processed using a 3D Fourier transform, which achieves high levels of sparsity due to the similarity of the image patches. Similarly, each 2D patch in the cluster can be vectorized and arranged into a matrix, which has low-rank properties, i.e., it can be represented with minimal error using a small number of singular values and vectors. Both of these properties are low-dimensional, since if a high-dimensional dataset has one of these properties, it can be approximated with high fidelity, using parameters (either transform domain coefficients or singular values and vectors) that have a total number much less than the dimensionality. Examples of such clusters and their low-dimensional properties are depicted in Fig. 1 for a slice containing the right coronary artery (RCA). Thus, each similarity cluster is highly likely to fit to a low-dimensional signal model, which allows artifacts due to undersampling and noise to be represented as perturbations to the low-dimensional structure, facilitating further processing, and preserving information locally. In the rest of the algorithm, these clusters are nonlinearly thresholded based on their low-dimensional properties to remove the perturbations caused by aliasing and artifacts. In the following sections, we describe the details of the LOST algorithm.

Learning the Anatomical Structure Using Similarity Clusters

To build clusters of similar anatomical content, we use a block matching algorithm that has also been used in video compression standards such as MPEG (26) and in image denoising (27). For each voxel in the image, we consider the 2D image block of size $N_b \times N_b$ (in the x - y direction), whose top left corner is located at that voxel location. This image block, denoted X_{ref} , is then compared with other $N_b \times N_b$ image blocks based on a normalized l_2 distance measure, $d(X_{\text{ref}}, X_{\text{other}}) = \|X_{\text{ref}} - X_{\text{other}}\|_2^2 / \|X_{\text{ref}}\|_2^2$. If this distance is less than a specified threshold, λ_{match} , then the two blocks are declared similar, and X_{other} is added to the similarity cluster of X_{ref} . To reduce the complexity, for any voxel, we only compare with 2D blocks within a specified search radius (N_{search} in x - y and N_{depth} in z direction), and limit the size of each cluster to be at most N_{cluster} . If more image blocks are matched to the reference block, only the N_{cluster} blocks with the highest degree of similarity are considered in the cluster for further processing.




FIG. 1. Coronary MRI contain areas of similar signal characteristics at various spatial locations. These similarity clusters can be treated as a three-dimensional (3D) object that exhibits transform domain sparsity or as a two-dimensional (2D) matrix with low-rank property. When dealing with undersampled acquisitions, the center of k -space is used to generate a low-resolution image from which these similarity clusters can be approximately identified and used in removing aliasing. The clusters are updated once better higher resolution estimates become available (FFT: fast Fourier transform, SVD: singular value decomposition).

Low-Dimensional Properties of Similarity Clusters

The block matching process ensures that the signal characteristics of the blocks in a given cluster are highly similar. This similarity enables the enhanced use of the low-dimensional properties for each cluster: transform-domain sparsity and low-rankedness. In this work, we concentrate on the transform-domain sparsity due to its modest computational requirements. How low-rankedness may be used to extend our approach is presented in Appendix A.

For efficient utilization of transform-domain sparsity, the 2D blocks in a similarity cluster are stacked into a 3D structure, and a 3D fast Fourier transform (FFT) is applied. The 2D FFT of each block promotes sparsity, and similar block-based transforms have been previously used in the JPEG standard (28). The transform along the third dimension of the similarity cluster further enhances sparsity since all the 2D blocks are chosen to have similar image content. For example, FFT in the third dimension would only keep the zero-frequency DC component if the blocks are equivalent.

If the similarity clusters were learned from a fully sampled high-resolution image, these properties would hold for each cluster. However, when dealing with undersampled k -space data, we only have access to low-resolution images from central k -space or estimates of the high-resolution image, which results in imperfections

in the identification of the similarity clusters. Furthermore, aliasing artifacts and noise will also be observed. Thus, we will deal with clusters, whose sparse components are corrupted by artifacts and noise. Therefore, nonlinear shrinkage has to be applied to each cluster to extract the desired low-dimensional structure of the clusters.

Nonlinear Shrinkage of Similarity Clusters

In the LOST de-aliasing strategy, the transform-domain sparsity of the similarity clusters is used to threshold noise and aliasing artifacts (Fig. 2). In our approach, we use shrinkage operators based on hard thresholding (capturing l_0 norm) in early iterations to reduce aliasing artifacts, and subsequently Wiener filtering (weighted l_2 norm) to reduce blurring artifacts, similar to image denoising algorithms (27).

A 3D FFT is applied to the similarity cluster of each voxel (x_i, y_j, z_l) , denoted by Λ_l . In the initial hard-thresholding steps, the FFT coefficients, $F_l = \mathbf{F}(\Lambda_l)$, are set to zero if their magnitude is below a threshold, τ_{ht} , and are unchanged if their magnitude is above the threshold. After a number of hard-thresholding iterations, Wiener filtering is applied to reduce the apparent blurring artifacts. In Wiener filtering, the filter coefficients for the (i, j, k) location in the Fourier domain are generated from the current estimate by




FIG. 2. The proposed LOST shrinkage uses the transform sparsity of similarity clusters generated. The sparsity property is used to remove aliasing by the application of a hard-thresholding operator in early iterations, and a Wiener filter after the estimates become more accurate.

$$W_l(i, j, k) = \frac{|F_l(i, j, k)|^2}{|F_l(i, j, k)|^2 + \tau_{\text{wie}}^2}, \quad [2]$$

where τ_{wie} is the Wiener filtering threshold. These filter coefficients and the FFT coefficients, F_l are then multiplied pointwise to generate the thresholded FFT coefficients. The thresholded FFT coefficients are subsequently transformed back to the image domain via 3D inverse FFT (IFFT).

At each iteration, the nonlinear filtering process is repeated for every similarity cluster, and since the clusters are not necessarily disjoint, there will be multiple thresholded estimates for 2D image blocks. To generate the final image estimate, individual block estimates are combined using weighted averaging (27). The intuition behind such a weighting scheme is to give smaller weights to noisier and more aliased blocks. The weights are generated, as explained below, for each cluster, and the corresponding blocks from that cluster are weighted with these weights in the final averaging stage. If hard thresholding is used, the weights are generated based on the number of non-zero coefficients remaining after thresholding, which will be denoted by N_l . In this case, the weights are given by

$$w_l^{\text{ht}} = \begin{cases} 1/N_l & \text{if } N_l \geq 1 \\ 1 & \text{if } N_l = 0 \end{cases}. \quad [3]$$

For Wiener filtering, the weights are generated using

$$w_l^{\text{wie}} = 1 / \sum_{i,j,k} |W_l(i, j, k)|^2. \quad [4]$$

After normalization, these weights are used to generate a new estimate for each $N_b \times N_b$ image block. The 2D image blocks themselves, which are also overlapping, are combined by simple averaging to generate the new image estimate (Appendix B). A Kaiser window is applied to each processed 2D block to reduce boundary artifacts.

METHODS

Reconstruction Algorithm: Implementation Details

The proposed reconstruction algorithm was implemented in two stages. Initially, low-resolution images were generated from the fully acquired center of k -space, using a Hanning window. These low-resolution images were used to learn the similarity clusters for the first stage. The parameters used for this process are $N_b = 8$, $N_{\text{search}} = 8$, $N_{\text{depth}} = 1$, $N_{\text{cluster}} = 16$, and $\lambda_{\text{match}} = 0.1$. Then for each coil, the iterative algorithm is applied (initialized to all-zero image), using hard-thresholding method for 25 iterations, generating $\hat{\mathbf{m}}_k^1$. The thresholding parameter for each coil, τ_{ht} , is chosen to be 0.05 times the largest (in absolute value) coefficient of the zero-filled coil image. This process results in removing the aliasing but results in certain blocking artifacts. In the second stage, the similarity clusters are updated again using $\hat{\mathbf{m}}_k^1$, with parameters $N_b = 4$, $N_{\text{search}} = 8$, $N_{\text{depth}} = 1$, $N_{\text{cluster}} = 16$ and $\lambda_{\text{match}} = 0.05$. The iterative algorithm is applied coil-by-coil, with initialization $\hat{\mathbf{m}}_k^1$, for a total of 15 iterations alternating between Wiener filtering and hard thresholding, generating the final estimates $\hat{\mathbf{m}}_k$. τ_{ht} and τ_{wie} are chosen to be 0.03 times the largest (in absolute value) coefficient of the previous coil estimate $\hat{\mathbf{m}}_k^1$.

The proposed method was implemented in Matlab (v7.6, MathWorks, Natick, MA), with the learning and nonlinear shrinkage portions implemented in C++. An implementation of LOST will be provided online.

In Vivo Imaging

All imaging sequences were implemented on a 1.5-T Philips Achieva (Philips Healthcare, Best, The Netherlands) system with a 5-channel cardiac-phased array receiver coil. The imaging protocol was approved by our institutional review board, and written informed consent was obtained from all participants.

A free-breathing 3D electrocardiographically gated steady state free precession (SSFP) sequence (echo time/pulse sequence repetition time = 2.1 ms/4.3 ms; field of view, $270 \times 270 \times 30 \text{ mm}^3$; flip angle, 90° spatial resolution, $1 \times 1 \times 3 \text{ mm}^3$) was used to image the right and left coronary arteries (29). A 2D right hemidiaphragm pencil-beam respiratory navigator with acceptance window of 5 mm was used for respiratory motion gating and tracking.

Coronary MRI: Restrospective Study

In two separate studies, right and left coronary MRI were acquired in two different subject cohorts. For each anatomy, 10 healthy adult subjects (four men, 22.0 ± 2.1 years for right, and four men, 30.7 ± 19.0 years for left) without contraindications to MR imaging were included. The k -space data were fully acquired, and were then exported and transferred to a standalone workstation to allow a retrospectively undersampled study. The data were undersampled by factors of 2, 3, and 4, by keeping the center 50×5 k_y - k_z -space lines, and randomly discarding outer k -space lines based on a zero-mean Gaussian distribution.

As a comparison to the proposed LOST strategy, the datasets were also reconstructed using l_1 minimization and TV regularization. l_1 Minimization was implemented

coil-by-coil via basis pursuit using the SPGL1 solver for large-scale l_1 regularized least squares problems (30). The solver was run for 300 iterations and orthogonal Daubechies-4 wavelets (15,21) were used as the sparsifying transform. TV regularization was implemented coil-by-coil via RecPF algorithm (31) using three-dimensional finite differences for $\Phi(\Psi\mathbf{m}) = \text{TV}(\mathbf{m})$, using 100 iterations and $\tau_{\text{TV}} = 0.0001$ (14,31). The final estimates were generated by root-sum-squares of the coil estimates for all reconstruction methods and for the reference image.

Coronary MRI: Prospective Study

Coronary MRI datasets were prospectively acquired at acceleration rates of 2, 3, and 4 in a total of four healthy subjects (two men, 22.0 ± 1.4 years). The imaging pulse sequence was modified to accommodate the random undersampling patterns and to mitigate effects due to eddy currents, gradient switching, and flow. Based on the undersampling factor, an undersampling pattern was generated by keeping the center 40×5 k_y - k_z lines and discarding edges randomly. The undersampling pattern was then stored as a lookup table. Before the acquisition, a phase reordering was performed, where the selected samples were sorted based on their k_y and k_z location in a radial fashion.

The prospectively undersampled images were reconstructed using both LOST, and l_1 minimization in the wavelet domain. Additionally, data were prospectively acquired with SENSE for the same acceleration rates, using uniform undersampling (8). The reconstructions made by the scanner were saved and used for comparison. Fully sampled images were also acquired for comparison. The nominal scan times for these exams were 3:11 min for fully sampled acquisition: 1:36 min, 1:04 min, and 49 s for acceleration rates of 2, 3, and 4, respectively, assuming 100% navigator efficiency and a heart rate of 70 beats/min.

Image and Statistical Analysis

For our retrospective coronary MRI study, both subjective and objective image analyses were performed to evaluate the three reconstruction methods (LOST, l_1 minimization and TV regularization) at different rates. For both left and right coronaries, a total of 10 imaging datasets were generated for each subject, including one from a fully sampled k -space, and three per method for acceleration rates of 2, 3, and 4. A qualitative assessment of coronary artery image quality was performed by an experienced independent blind reader with >10 years coronary MRI experience using a four-point scale system (1 = poor, 4 = excellent) (32). Separate scores were given for the proximal, mid, and distal segments of the RCA; and left main (LM), proximal, mid left anterior descending artery (LAD), and left circumflex artery (LCX). These scores were combined for the right and left systems, respectively, before statistical analysis.

The Soap bubble (33) tool was used to quantitatively evaluate the RCA vessel sharpness. Vessel sharpness scores were calculated for both sides of the vessel. Final normalized sharpness was defined as the average score of both sides normalized by the center of vessel intensity. For quantitative measurement, the mean square

error (MSE) of each reconstruction was also calculated. The normalized MSE was then calculated by dividing each individual MSE by the sum of the squared l_2 norms of the reference coil images.

Imaging scores, sharpness, and MSE are presented as mean $\pm$ one standard deviation. The signed rank test was used for imaging scores to test for the null hypothesis that the central tendency of the difference was zero at different acceleration rates. All statistical analyses were performed using SAS (v9.2, SAS Institute, Cary, NC). The normalized sharpness scores and normalized MSE were compared using the paired t test. A P value of <0.05 was considered to be significant.

RESULTS

Figure 3 shows a sample 2D slice from a 3D coronary MRI dataset from fully sampled k -space data, and reconstructed images using LOST at acceleration rates of 2, 3, and 4, with corresponding reconstructions using l_1 minimization in the wavelet domain and TV regularization in middle and bottom row, respectively. With acceleration greater than 2, l_1 minimization and TV regularization suffer from blurring and artifacts. In all methods, there is a degradation of image quality as the acceleration rate is increased, although the degradation of LOST images is less severe. Figure 4 depicts reformatted images of LAD, LCX, and proximal RCA from a 3D left coronary MRI dataset reconstructed using the three approaches. Both l_1 norm minimization and TV regularization suffer from blurring at rates higher than 2. LOST allows for improved visualization at higher rates but also has artifacts.

Figure 5 summarizes the qualitative assessment of the coronaries, which indicate LOST outperforms both l_1 norm minimization in the wavelet domain and TV regularization at all rates. Statistical analysis shows that the difference between LOST and the other two methods are significantly different at all rates ($P < 0.05$). There is a significant difference between TV regularization and l_1 minimization at rate 4 for the RCA, in favor of TV regularization. There is also a significant difference among TV regularization and l_1 minimization at rate 2 for LAD/LM/LCX.

Table 1 summarizes normalized sharpness and MSE for RCA and normalized MSE for LAD/LM/LCX, using LOST, l_1 norm minimization in the wavelet domain and TV regularization. The mean value of MSE for LOST was lower than l_1 norm minimization and TV regularization. Except for acceleration rate 2 of the LAD/LM/LCX, TV regularization consistently had a lower MSE than l_1 norm minimization in the wavelet domain, and their difference at rate 2 was not statistically significant. When the three methods were compared pairwise, there were statistically significant differences between all methods, except for rate 2 between TV and l_1 norm minimization.

LOST had the highest RCA sharpness score among all the reconstruction methods. l_1 norm minimization in the wavelet domain had a higher sharpness score compared with TV regularization, which is consistent with the qualitative assessment. When compared with the fully sampled reference, there were no significant differences for LOST at rate 2. All the other reconstructions exhibited significant differences with respect to the reference ($P < 0.05$).




FIG. 3. Retrospective accelerated RCA reconstruction: An example slice of a 3D coronary MR image from fully sampled data, LOST (top row) and l_1 norm minimization in wavelet domain (middle row) and total variation (TV) regularization (bottom row) at rates of 2, 3, and 4. RCA is readily visualized at all rates for images using LOST. The RCA is highly blurred for rate-4 l_1 minimization and TV regularization. Furthermore, more apparent artifacts and blurring are observed for these two techniques at rates 3 and 4 compared with LOST (AO: aortic root; RV: right ventricle; LV: left ventricle; RCA: right coronary artery).

Reformatted RCA images from a prospective acquisition are depicted in Fig. 6 from fully sampled acquisition, LOST, and l_1 norm minimization with random prospective undersampling for rates 2, 3, and 4. The SENSE reconstruction with a uniform undersampling pattern for the same subject is also shown. At rate 2, the three reconstructions provide images where the RCA is clearly visible, although l_1 norm minimization suffers from blurring, and SENSE has noise amplification. The differences become more pronounced at higher rates of 3 and 4. Figure 7 depicts a reformatted image from a prospective LAD/LM/LCX acquisition, which shows similar image quality to RCA. At rates 3 and 4, LOST provides visualization of the arteries, whereas the proximal RCA and LCX are no longer visible in SENSE, and all the arteries are blurred for l_1 norm minimization. There are inter-exam variations among images acquired using uniform and random undersampling, as well as the fully sampled scan, due to variations in the subjects' breathing pattern.

DISCUSSION

We have developed an improved reconstruction technique for undersampled k -space data, which exploits the image structure to remove aliasing and imaging artifacts, and

evaluated this novel method in coronary MRI. The use of similarity clusters enables high levels of sparsity, because the blocks within the clusters exhibit similar signal characteristics. This is stronger than merely assuming sparsity of signals in a transform-domain or assuming a probabilistic model of sparsity structure in images, as conventional CS reconstructions do. The learning procedure adaptively determines the structure of the image, thus the variations between unique anatomies, patients, imaging parameters, and contrast level are captured for the image of interest, although further investigation about its utility with these variations is necessary. Furthermore, the learning procedure also identifies variations in different coil images, even when the method is used coil-by-coil.

The shortcomings of fixed transforms in representing different anatomies were previously reported (34–36), and transforms fine tuned to the given anatomy using a dictionary learning algorithm were used (24). Dictionary learning methods are based on a global optimization model, where a dictionary is generated for all the components of the image (with the components being the parameter direction (34), patches (35), or columns of the image (36)). Thus, a dictionary is generated with the goal to simultaneously sparsify all the components, and each




FIG. 4. Retrospective accelerated left coronary MRI: Reformatted axial images of the RCA, left anterior descending artery (LAD), and left circumflex artery (LCX) from fully sampled 3D left coronary MRI data (reference), and from LOST (top row), l_1 minimization (middle row), and TV regularization (bottom row) reconstructions for accelerated acquisition rates of 2, 3, and 4. The coronaries can be readily visualized in all images, although blurring is apparent at rates 3 and 4 for l_1 minimization and TV regularization. (LAD: left anterior descending; LCX: left circumflex artery).




FIG. 5. Mean image reading scores and standard deviation for the right and left coronaries, for all reconstructions at rates 2, 3, and 4. LOST outperforms the other methods at all rates. The significant differences ($P < 0.05$) are marked with *.

Table 1
Normalized Sharpness and Normalized MSE for Right Coronary Acquisitions, and Normalized MSE for Left Coronary Acquisitions, Using LOST, l_1 Minimization and TV Regularization

	REF	LOST dealiasing				l_1 Minimization				TV Regularization			
		R2		R3		R2		R3		R2		R3	
		R4		R4		R4		R4		R4		R4	
Right	MSE (norm.)	0.06 ± 0.03	0.12 ± 0.04	0.15 ± 0.05	0.11 ± 0.03	0.19 ± 0.05	0.24 ± 0.06	0.09 ± 0.03	0.13 ± 0.04	0.15 ± 0.05			
	Sharp (norm.)	0.54 ± 0.09	0.52 ± 0.08	0.50 ± 0.08	0.52 ± 0.08	0.47 ± 0.09	0.42 ± 0.09	0.49 ± 0.09	0.46 ± 0.08	0.41 ± 0.08			
Left	MSE (norm.)	0.05 ± 0.02	0.10 ± 0.04	0.13 ± 0.04	0.10 ± 0.03	0.17 ± 0.05	0.21 ± 0.06	0.11 ± 0.03	0.12 ± 0.04	0.14 ± 0.05			

component is thresholded individually such that its representation has at most a predefined number of nonzero coefficients. In contrast, our model uses clusters that only have blocks of similar content, and each block can be part of multiple clusters, which reduces the artifacts associated with standard thresholding methods (27). Furthermore, as dictionary learning methods try to simultaneously sparsify the components based on an l_2 constraint, the relevant anatomical features are not necessarily captured in an equally sparse manner, which may lead to artifacts. The dictionary learning stage can be performed using training data (34,35), which might not reflect the variability, contrast, and image features among patients. Alternatively, the dictionary and the image can be approximated together, where the dictionary is updated with least squares constraints, based on the current image estimate, which is the solution of a CS reconstruction problem that uses the previous estimate for the nonparametric dictionary (36). This approach suffers from high-computational burden and the potential to converge to a local minimum, because the dictionary update only relies on a least squares solution without further constraints, although the measurements are highly incomplete.

In our implementation, we used different shrinkage techniques depending on the iteration number. At the earlier iterations, hard thresholding is preferred as it effectively removes aliasing. In our experiments, we found that hard thresholding is much more effective than soft thresholding (capturing l_1 norm of each cluster) for removing aliasing in similarity clusters. Wiener filtering is not applicable in earlier iterations, as the aliasing artifacts produce unreliable filter coefficients. However, in later iterations, Wiener filtering with appropriate filter coefficients is preferred because of its effectiveness in removing blocking and noise artifacts. These observations are consistent with the results in image denoising literature (27).

In terms of computational complexity, the complexity of the data consistency stage is dominated by an FFT and an IFFT, both of $O(N \cdot \log(N))$, where N is the image size; which is standard in many reconstruction algorithms. In the dealiasing stage, the shrinkage operations scale linearly with the dimensions. For hard thresholding and Wiener filtering, the complexity is $O(N_b^2 \cdot N_{\text{cluster}} \cdot \log(N_b \cdot N_{\text{cluster}}))$ per voxel. In our implementation, we limit ourselves to small constant block and cluster sizes, which results in a run-time not more than wavelet-based thresholding. The block-matching portion before the iterations, requires pairwise comparisons for each block within the search radius, resulting in a computational complexity of $O((N_b \cdot N_{\text{search}})^2 \cdot N_{\text{depth}})$, which takes almost as long as the iterative portion itself. Further accelerations in the learning process can benefit the overall algorithm, enabling more frequent update of similarity clusters. For instance, rather than sliding through each voxel one-by-one as we currently do in our algorithm, several voxels might be skipped at once (27). The run-times for the algorithms were 7 min for TV (RecPF), 98 min for LOST, and 385 min for l_1 norm minimization (SPGL1) for reconstructing five coil images, each of size $544 \times 270 \times 19$, on a standard Lenovo workstation with a 2.66-GHz central processing unit and 8-GB RAM. We also note that since most of the processing in our




FIG. 6. Prospective accelerated RCA acquisition: Reformatted images from a prospectively undersampled targeted RCA acquisition, from fully sampled data, LOST (top row), l_1 minimization (middle row), and SENSE (bottom row) for acquisition rates of 2, 3, and 4. Random undersampling was utilized for the acquisitions reconstructed with LOST and l_1 minimization, uniform undersampling was utilized for SENSE. At rate 2, LOST reconstruction is visually very similar to the reference, whereas l_1 minimization has blurring, and SENSE has folding artifacts. At higher rates, LOST still manages to provide improved visibility, while blurring and noise amplification are apparent in l_1 minimization and SENSE images, respectively.

algorithm is done independently for each voxel, as opposed to a linear transformation across the image, a parallel implementation using a graphics processing unit (GPU) may result in significant run-time gains, bringing the running time to more clinically feasible levels.

Some techniques were used in our implementation to simultaneously meet requirements of both speed and quality of reconstruction, but alternatives may be considered as well. For instance, we used an l_2 distance-based similarity measure, which is advantageous because the triangle inequality guarantees that the maximum distance between any two elements in a given cluster will be bounded. However, other more computationally demanding similarity measures such as mutual information (37) and shape-adaptive block matching (38) or 3D block matching may be used.

We have used FFT as our 3D transform for the clusters. The use of FFT, especially in the third dimension (along the blocks) is crucial for ensuring high levels of transform domain sparsity. We note that as the similarity clusters contain only blocks of similar signal con-

tent, the periodic boundary conditions of FFT along the third dimension do not hinder the representation sparsity. The transform in the other two dimensions (within the blocks) is less critical, but we used FFT based on the availability of fast C++ implementations and for its natural capacity to handle complex data. Further investigation is required for optimization. We also note that due to the contents of similarity clusters, a 3D FFT not only captures the low-resolution characteristics but most of the characteristics. This in conjunction with the use of a fixed thresholding value (and not a fixed level of sparsity) for all clusters ensures that high-resolution features in areas with substantial signal content are not lost, as confirmed by the quantitative sharpness measurements.

In the highly overcomplete representation used in this work, we allow for overlapping neighborhoods and for image patches to be represented in multiple different ways in different clusters. Thus, a thresholding based on a fixed number of coefficients for each cluster is not appropriate in general. Finally, the processed patches are




FIG. 7. Prospective accelerated left coronary MRI: Reformatted axial images depicting RCA, LCX, and LAD from a prospectively undersampled targeted left coronary MRI acquisition, from fully sampled data, LOST (top row), l_1 minimization (middle row), and SENSE (bottom row) for acquisition rates of 2, 3, and 4. At rate 2, the three methods reconstruct images of similar quality, although l_1 minimization exhibits blurring. At rates 3 and 4, l_1 minimization and SENSE suffer from blurring and noise amplification, respectively, and the coronaries can no longer be defined.

combined in a nonlinear fashion based on the effects of thresholding for that cluster. Hence, the method does not really fit the traditional CS framework, and the compressibility analysis of (15) cannot be applied in a straightforward manner.

The size of the fully sampled central k -space for this work was determined empirically to ensure that the low-resolution image generated from these k -space lines did not exhibit ringing artifacts. This ensures that the algorithm does not learn these artifacts as part of the image. If a smaller part of the central k -space is acquired, an estimate can be generated using standard CS techniques, and this estimate can be used to determine the similarity clusters.

The proposed dealiasing approach may also be used in combination for combining parallel imaging and CS (20–23), as well as with non-Cartesian trajectories. Although this study did not test the optimal way of combining the proposed method with parallel imaging techniques, the local processing of the blocks may yield more robust reconstruction to imperfect coil sensitivity maps. The use of a phased-array coil with a higher number of channels may yield higher acceleration rates, because images will be better localized in different coils. The proposed method may also improve image reconstruction in other anatomical or functional imaging, which requires further investigation.

CONCLUSIONS

We have developed an improved reconstruction technique for undersampled acquisitions that learns and utilizes the structure of images being reconstructed and evaluated its performance in coronary MRI.

APPENDIX A

Extensions: Low Rank-Based Processing of Similarity Clusters

Low-rank properties of matrices have been an active area of recent research, with applications in low-rank matrix completion from incomplete measurements and robust principal component analysis techniques (39,40). To utilize low rankedness of clusters, each 2D patch in the cluster is rearranged into an $N_b^2 \times 1$ vector, and these vectors are treated as the columns of a matrix describing the cluster. Because the data in each 2D block has similar anatomical features, the matrix can be well approximated by a low-rank matrix. As a simple example, if all the 2D blocks had the same signal content, then the matrix would have rank 1.

To utilize the low-rank properties of the cluster Λ_j , we first generate the cluster matrix $\mathbf{M}_{\Lambda_j}$ by rearranging the

2D blocks in Λ_l to column vectors. By the low-rank properties of the similarity clusters, $\mathbf{M}_{\Lambda_l}$ can be well approximated by a low-rank matrix, $\mathbf{L}_{\Lambda_l}$, corrupted by artifacts and noise. Classical principal component analysis seeks to find the best rank- k matrix that approximates $\mathbf{M}_{\Lambda_l}$ in the l_2 sense. Although this is optimal in the presence of Gaussian noise, principal component analysis is not robust to gross and non-Gaussian errors caused by artifacts and noise, and a more robust version of principal component analysis, denoted principal component pursuit (PCP) has recently been proposed in (40). Here, the noise and artifacts are modeled as a sum of gross errors, $\mathbf{E}_{\Lambda_l}$ and small errors $\mathbf{Z}_{\Lambda_l}$, resulting in the model

$$\mathbf{M}_{\Lambda_l} = \mathbf{L}_{\Lambda_l} + \mathbf{E}_{\Lambda_l} + \mathbf{Z}_{\Lambda_l}. \quad [\text{A1}]$$

The low rank matrix can then be recovered using a convex relaxation of rank minimization

$$\min_{\mathbf{L}, \mathbf{E}} \|\mathbf{L}\|_* + \lambda_{\text{pcp}} \|\mathbf{E}\|_1 + \mu_{\text{pcp}} \|\mathbf{M}_{\Lambda_l} - (\mathbf{L} + \mathbf{E})\|_F^2, \quad [\text{A2}]$$

where $\|\cdot\|_*$ denotes the nuclear norm of the matrix (sum of the absolute values of the singular values), $\|\cdot\|_1$ denotes the l_1 norm of the matrix (sum of the absolute values of the entries), and $\|\cdot\|_F^2$ denotes the squared Frobenius norm of the matrix (sum of the magnitude-squares of the entries), and λ_{pcp} and μ_{pcp} are optimization weights (40). This minimization can be solved using iterative approaches (40). However, this results in a complexity of $O(\max(N_b^2, N_{\text{cluster}})^6)$ per voxel, which prohibits a comprehensive study of this implementation. We note that a GPU implementation may enable the use of this technique, which is attractive because it can handle non-Gaussian artifacts.

APPENDIX B

Weighted Averaging of Processed Blocks

The weighted averaging approach used to generate a final image estimate from thresholded blocks in similarity clusters can be formally described as

$$\mathbf{m}^{(t+1)} = \sum_{(x_0, y_0, z_0)} \frac{\sum_{l: (x_0, y_0, z_0) \in \Lambda_l} w_l^{\text{filt}} \chi_{x_0, y_0, z_0}(\mathbf{F}^{-1}(\Phi_{\text{filt}}(\mathbf{F}(\Lambda_l))))}{\sum_{l: (x_0, y_0, z_0) \in \Lambda_l} w_l^{\text{filt}}} \bigg/ \sum_{(x_0, y_0, z_0)} \mathbf{I}(x_0, y_0, z_0), \quad [\text{B1}]$$

where Φ_{filt} specifies which nonlinear filter has been used (ht or wie), w_l^{filt} specifies the corresponding weight for the appropriate cluster. $\chi_{x_0, y_0, z_0}(\cdot)$ is a function that keeps the 2D image block whose top left corner is located at (x_0, y_0, z_0) among the blocks in the cluster $(\cdot)$ and maps this to its appropriate location in the image, and $\mathbf{I}(x_0, y_0, z_0)$ is the indicator function that is 1 for

$$\{(x, y, z) : x_0 \leq x \leq x_0 + N_b - 1, y_0 \leq y \leq y_0 + N_b - 1, z = z_0\} \quad [\text{B2}]$$

and 0 otherwise.

REFERENCES

- McGibney G, Smith MR, Nichols ST, Crawley A. Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI. *Magn Reson Med* 1993;30:51–59.
- Meyer CH, Hu BS, Nishimura DG, Macovski A. Fast spiral coronary artery imaging. *Magn Reson Med* 1992;28:202–213.
- Peters DC, Korosec FR, Grist TM, Block WF, Holden JE, Vigen KK, Mistretta CA. Undersampled projection reconstruction applied to MR angiography. *Magn Reson Med* 2000;43:91–101.
- Barger AV, Block WF, Toropov Y, Grist TM, Mistretta CA. Time-resolved contrast-enhanced imaging with isotropic resolution and broad coverage using an undersampled 3D projection trajectory. *Magn Reson Med* 2002;48:297–305.
- Mistretta CA, Wieben O, Velikina J, Block W, Perry J, Wu Y, Johnson K. Highly constrained backprojection for time-resolved MRI. *Magn Reson Med* 2006;55:30–40.
- Winkelmann S, Schaeffter T, Koehler T, Eggers H, Doessel O. An optimal radial profile order based on the golden ratio for time-resolved MRI. *IEEE Trans Med Imaging* 2007;26:68–76.
- Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. *Magn Reson Med* 1997;38:591–603.
- Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. *Magn Reson Med* 1999;42:952–962.
- Griswold MA, Jakob PM, Heidemann RM, Nittka M, Jellus V, Wang J, Kiefer B, Haase A. Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* 2002;47:1202–1210.
- Prieto C, Mir R, Batchelor PG, Hill DLG, Guarini M, Irarrazaval P. Reconstruction of under-sampled dynamic images by modeling the motion of object elements. *Proceedings of the 14th Scientific Meeting of ISMRM, Seattle; 2006. p 687.*
- Sharif B, Derbyshire J, Faranesh A, Lederman R, Bresler Y. Real-time non-gated cardiac MRI using PARADISE: doubly adaptive accelerated imaging. *Proceedings of the 17th Scientific Meeting of ISMRM, Honolulu; 2009. p 767.*
- Odille F, Uribe S, Batchelor PG, Prieto C, Schaeffter T, Atkinson D. Model-based reconstruction for cardiac cine MRI without ECG or breath holding. *Magn Reson Med* 2010;63:1247–1257.
- Tsao J, Boesiger P, Pruessmann KP. k-t BLAST and k-t SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations. *Magn Reson Med* 2003;50:1031–1042.
- Block KT, Uecker M, Frahm J. Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint. *Magn Reson Med* 2007;57:1086–1098.
- Lustig M, Donoho DL, Pauly JM. Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn Reson Med* 2007;58:1182–1195.
- Vasanawala SS, Alley MT, Hargreaves BA, Barth RA, Pauly JM, Lustig M. Improved pediatric MR imaging with compressed sensing. *Radiology* 2010;256:607–616.
- Donoho DL. Compressed sensing. *IEEE Trans Inf Theory* 2006;52:1289–1306.
- Candès EJ, Romberg J, Tao T. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans Inf Theory* 2006;52:489–509.
- Gamper U, Boesiger P, Kozerke S. Compressed sensing in dynamic MRI. *Magn Reson Med* 2008;59:365–373.
- Otao R, Kim D, Axel L, Sodickson DK. Combination of compressed sensing and parallel imaging for highly accelerated first-pass cardiac perfusion MRI. *Magn Reson Med* 2010;64:767–776.
- Lustig M, Pauly JM. SPIRiT: iterative self-consistent parallel imaging reconstruction from arbitrary k-space. *Magn Reson Med* 2010;64:457–471.
- Fischer A, Seiberlich N, Blaimer M, Jakob P, Breuer F, Griswold M. A combination of nonconvex compressed sensing and GRAPPA (CS-GRAPPA). *Proceedings of the 17th Scientific Meeting of ISMRM, Honolulu; 2009. p 2813.*
- Liang D, Liu B, Wang J, Ying L. Accelerating SENSE using compressed sensing. *Magn Reson Med* 2009;62:1574–1584.
- Elad M, Aharon M. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Trans Image Process* 2006;15:3736–3745.
- Figueiredo M, Nowak R. An EM algorithm for wavelet-based image restoration. *IEEE Trans Image Process* 2003;12:906–916.

26. Hanzo L, Cherriman P, Streit J. Video compression and communications: from basics to H.261, H.263, H.264, MPEG4 for DVB and HSDPA-style adaptive turbo-transceivers: Wiley-IEEE Press; 2007.
27. Dabov K, Foi A, Katkovnik V, Egiazarian K. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Trans Image Process* 2007;16:2080–2095.
28. Donoho DL, Vetterli M, DeVore RA, Daubechies I. Data compression and harmonic analysis. *IEEE Trans Inf Theory* 1998;44:2435–2476.
29. Stuber M, Botnar RM, Danias PG, Sodickson DK, Kissinger KV, Van Cauteren M, De Becker J, Manning WJ. Double-oblique free-breathing high resolution three-dimensional coronary magnetic resonance angiography. *J Am Coll Cardiol* 1999;34:524–531.
30. van den Berg E, Friedlander MP. Probing the Pareto frontier for basis pursuit solutions. *SIAM J Sci Comput* 2008;31:890–912.
31. Yang J, Zhang Y, Yin W. A fast alternating direction method for TVL1-L2 signal reconstruction from partial Fourier data. *IEEE J Sel Top Signal Process* 2010;4:288–297.
32. Kim WY, Danias PG, Stuber M, Flamm SD, Plein S, Nagel E, Langerak SE, Weber OM, Pedersen EM, Schmidt M, Botnar RM, Manning WJ. Coronary magnetic resonance angiography for the detection of coronary stenoses. *N Engl J Med* 2001;345:1863–1869.
33. Etienne A, Botnar RM, Van Muiswinkel AM, Boesiger P, Manning WJ, Stuber M. “Soap-bubble” visualization and quantitative analysis of 3D coronary magnetic resonance angiograms. *Magn Reson Med* 2002;48:658–666.
34. Doneva M, Bornert P, Eggers H, Stehning C, Senegas J, Mertins A. Compressed sensing reconstruction for magnetic resonance parameter mapping. *Magn Reson Med* 2010;64:1114–1120.
35. Bilgin A, Kim Y, Liu F, Nadar MS. Dictionary design for compressed sensing MRI. *Proceedings of the 18th Scientific Meeting of ISMRM, Stockholm*; 2010. p 4887.
36. Otazo R, Sodickson D. Adaptive compressed sensing MRI. *Proceedings of the 18th Scientific Meeting of ISMRM, Stockholm*; 2010. p 4867.
37. Pluim JP, Maintz JB, Viergever MA. Mutual-information-based registration of medical images: a survey. *IEEE Trans Med Imaging* 2003;22:986–1004.
38. Foi A, Katkovnik V, Egiazarian K. Pointwise shape-adaptive DCT for high-quality denoising and deblocking of grayscale and color images. *IEEE Trans Image Process* 2007;16:1395–1411.
39. Candès EJ, Recht B. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics* 2009;9:717–772.
40. Zhou Z, Li X, Wright J, Candès EJ, Ma Y. Stable principal component pursuit. *Proceedings of the IEEE International Symposium on Information Theory, Austin*; 2010. p 1518–1522.