Compressed Sensing Reconstructions for Dynamic Contrast Enhanced MRI

Kevin T. Looby
klooby@stanford.edu

Abstract

The temporal resolution necessary for dynamic contrast enhanced (DCE) magnetic resonance imaging (MRI) requires scan accelerations of more than $8\times$ a typical scan time [1]. Aggressive undersampling during data acquisition can facilitate adequate temporal resolution at the cost of reduced spatial resolution. [1] presents a method for attaining the MR accelerations necessary while preserving spatial resolution via compressed sensing methods. This project seeks to implement this method with an additional compressed sensing component.

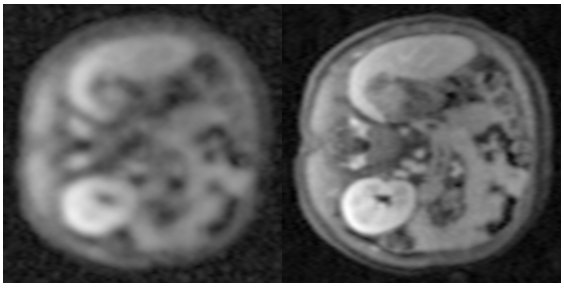

I. Background

A. Dynamic Contrast Enhanced MRI

Dynamic contrast enhanced magnetic resonance imaging allows for the characterization of lesions (*e.g.*, tumors, damaged tissue, etc.) as well as examination of the performance of the renal system. Visualization of the flow dynamics of contrast agents throughout tissues over time can be used for tissue evaluation and characterization based on the tissue uptake rates of the contrast agents.

Contrast agent uptake typically completes within a few minutes, about an order of magnitude faster than the scan rate of a conventional full-resolution MR acquisition. The problem is compounded in pediatric DCE MRI, as the flow dynamics occur over shorter time periods in very young patients. Parallel imaging acquisitions undersampled at a factor of $18\times$ can achieve a temporal resolution of roughly 3 seconds, which is sufficient for pediatric DCE MRI [1].

B. Parallel MRI Reconstructions with SPIRiT

SPIRiT (Iterative Self-consistent Parallel Imaging Reconstruction) is a parallel imaging reconstruction technique used to reconstruct multi-coil MR datasets. The algorithm uses a projection onto convex sets type approach similar to those to be detailed in sections II.c1 and II.c2 in order to utilize data redundancy from the multiple receiver channels to form an accurate reconstruction.

The details of this method are not the focus of this project and will not be discussed further, but as SPIRiT forms an integral component of my procedure, it is important that this basic explanation be given in advance. A MATLAB implementation of SPIRiT was provided courtesy of Michael Lustig and John Pauly.

II. Methods

Materials

All datasets used in this project were provided courtesy of Tao Zhang and John Pauly, both of Stanford's Magnetic Resonance Systems Research Lab (MRSRL). Two sets of two-dimensional pediatric abdominal cross-section dynamic contrast enhanced datasets were provided and a dynamic contrast enhanced image phantom were provided. MATLAB code for performing SPIRiT [2] reconstructions and wavelet transforms were made available by Michael Lustig and John Pauly.

Procedure Outline

A brief overview of the proposed data restoration and reconstruction method is provided. The following sections will elaborate further on each step.

1) Choose appropriate undersampling pattern
2) Acquire parallel DCE MR data with chosen sampling
3) Compress coils with eigenspace representation
4) Perform SPIRiT reconstruction with spatial and temporal compressed sensing data consistency enforcement+
5) Combine multi-coil data and form final reconstruction

A. Undersampling

The attainable results of the proposed method depend heavily on the choice of undersampling pattern. The denoising steps described later are all of projection on convex sets (POCS) type algorithms. As will be discussed in further detail, the undersampling pattern should be chosen in order to produce aliasing artifacts with characteristics similar to incoherent white noise. The POCS algorithms seek to iteratively reduce the noise levels while maintaining the power of the true signal. This is best achieved with low-power incoherent white noise.

By nature, aliasing artifacts produce image noise with inherent coherence. In the case of uniform undersampling, the aliasing artifacts produced are exact replicas of the original signal, and are therefore highly coherent. Because the aliasing replicas have similar power characteristics to the true signal, POCS algorithms will be unable to remove any noise. However, more strategically-chosen undersampling methods can produce lower noise floors. The following characteristics describe a sampling scheme that will produce aliasing with the desired incoherent white noise similarity:

- Randomness
- Variable density, particularly higher-density of samples in lower spatial frequency components
- Maximal coverage of k-space

Figure 1 shows a fully sampled DCE MR phantom reconstruction and its Fourier transform along with four k-space undersamplings and their multi-coil inverse Fourier transform reconstructions. The reconstructions were performed as the per-pixel square root of the sum of the cross-coil image intensity values. That is,

$$I[x,y] = \sqrt{\sum_c |\mathcal{F}^{-1}(k_c[x,y])|}, \qquad (1)$$

where I denotes the reconstructed image, c the set of coils, k_c the k-space data for each coil, and $\mathcal{F}^{-1}$ the inverse Fourier transform operator. For each sampling scheme, the phantom dataset has been undersampled by a factor of 16, meaning that only one sixteenth of the original data is kept.

A qualitative examination of the reconstruction from the variable-density Poisson disk undersampling shows that this sampling pattern achieves the desired characteristics. The variable-density Poisson disk sampling utilizes random sampling, with larger sample density at lower spatial frequencies. Additional k-space coverage is obtained by enforcing the additional constraint that no sampled point can be within a predefined distance of any other sampled point. This is the sampling method used in [1], and all reconstructions shown throughout the remainder of this paper will be performed on Poisson disk undersampled datasets.

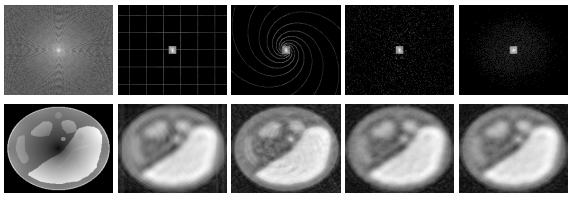

Fig. 1. k-space sampling schemes (top row, from left-to-right: fully-sampled, uniform undersampling, spiral undersampling, random uniform-density undersampling, variable-density Poisson disk undersampling) with corresponding multi-coil inverse Fourier transform reconstructions (bottom row). Coils were combined as the per-pixel square root of the sum of the cross-coil intensity values. Note: each k-space sampling scheme shown additionally includes the 16×16 block of lowest spatial frequencies. This block is referred to as the "calibration region" and is necessary for self-calibrating parallel MR reconstructions (e.g., SPIRiT) [2].

B. Parallel Imaging

Parallel imaging has long been used as a method for MRI acceleration. In a parallel MR acquisition, data is acquired on multiple receiver coils, each of which captures the volume with varying regional sensitivity. That is, each coil senses a particular region of the volume with high sensitivity but that sensitivity decreases with distance. The data from all of the coils is combined in order to produce a reconstruction over the entire volume.

By simultaneously acquiring data on multiple receiver coils, a division-of-labor is created by letting each receiver sample more sparsely. Utilizing redundant data samples from across the multiple coils, it is possible to create nearly perfect reconstructions with a proper coil data-combining algorithm. In this project, I make use of SPIRiT [2], using a MATLAB implementation courtesy of Michael Lustig (of Berkeley) and John Pauly.

At high acceleration factors, these algorithms fail, as the k-space data is too sparse. The rest of this project approaches using parallel imaging and compressed sensing to estimate and restore the unsampled spatial frequencies such that a sufficiently high resolution reconstruction becomes possible.

1) Coil Compression: The sparsity of the spatial frequency information can be dramatically reduced by transforming the data to different space. Because each channel measures the same volume, the data across the multiple channels is highly compressible. An eigenspace representation is an efficient transformation for parallel MR data.

In the dataset shown in figure 3 (top), the original dataset contained data from eight significant channels. The eigenvalue decomposition

of this dataset results in only four significant eigenvalues. In this "eigencoil" representation, the significant data has been compressed into fewer channels. The lower energy channels can be removed from the reconstruction in order to both denoise the data, as these channels contain little more than noise, and additionally reduce the size of the data needed for storage and processing.

For the phantom image dataset, my method retains the data corresponding to the four largest eigenvalues. Figure 2 shows the multicoil reconstruction (equation 1) of the data before compression (left) and the reconstruction using the four most-significant eigencoils. The pediatric DCE MRI datasets provided by John Pauly and Tao Zhang had been previously compressed to the eight highest-energy eigencoils.

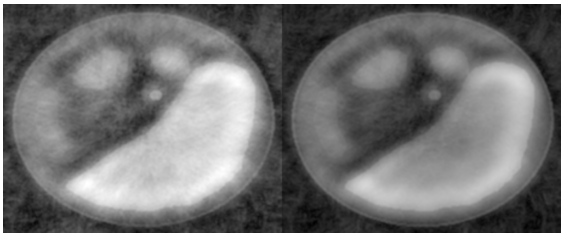

Fig. 2. Reconstruction without coil compression (left) and reconstruction of four highest-energy eigencoils.

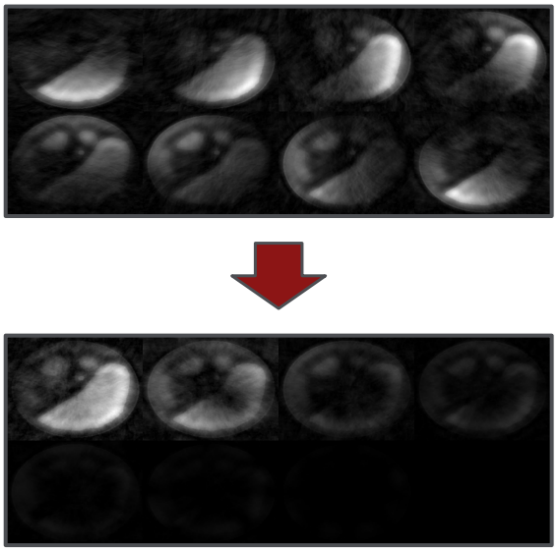

Fig. 3. Because the channels of a parallel MR dataset each sense the full field of view of the volume, there is a large amount of redundancy across the channels. This feature causes the eigenspace representation to be a very efficient compression.

C. Compressed Sensing in MRI

In general, compressed sensing refers to the process of measuring very few random linear combinations of signal values. The quantity of samples collected is well below the number needed to uniquely reconstruct the signal from the samples alone. Nonlinear methods can be used, however, to perform high-quality reconstructions.

In MRI, the undersampled k-space samples comprise the linear signal value combinations. Compressed sensing theory in this case indicates that it is possible to form an accurate reconstruction using a small sampling of k-space, rather than a full Cartesian grid [3]. As stated in section 2.a, it is crucial that the undersampling pattern be random, because random undersampling produces a viable compressed sensing denoising problem.

This problem can be posed as an l_1-penalized convex optimization problem. Presented in equation 2 is a slight modification to the problem posed by [4].

$$\text{argmin}\frac{1}{2}||S\hat{x} - y||_2^2 + \lambda|\hat{x}|_1 \tag{2}$$

Here, S indicates a sparsifying operator (*i.e.* a transform that produces a sparse representation of a signal), $\hat{x}$ is the estimated signal, y is the known signal information, *i.e.*, the acquired samples, and λ is a scalar value that enforces the l_1 norm penalty.

Solving this constrained problem attempts to denoise the undersampled image by computing a sparse representation of the image data while maintaining data consistency with the sampled spatial frequency information [3][4]. An iterative projection onto convex sets (POCS) type algorithm can be used to solve this problem for a given sparsifying transform, S. The following sections will detail the solution algorithm and sparsifying transforms chosen.

1) Spatial Compressed Sensing: Magnetic resonance images have an exploitable sparse representation in the wavelet domain [3][4][5]. Letting W represent a wavelet transform operator, equation 2 can be reformulated as

$$\text{argmin}\frac{1}{2}||W\hat{x} - y||_2^2 + \lambda|\hat{x}|_1. \tag{3}$$

A computational solution to this problem can be achieved through the following iterative POCS type algorithm [4].

1) Transform k-space data to image space with the inverse Fourier transform
2) Apply a wavelet transform to the image space data
3) Threshold low-energy wavelet coefficients
4) Apply the inverse wavelet transform to return to image space
5) Return to the k-space representation via the forward Fourier transform

6) Restore the data known to be true (the sampled k-space data)
7) Repeat 1 through 6 until sufficient convergence is achieved

Steps 1 through 4 are shown visually in figure 4. In my implementation, the threshold is originally set at 40% of the largest wavelet coefficient (only coefficients with energy no less than 40% of the energy of the largest coefficient are retained). I define sufficient convergence to be

$$1.0 > \frac{|\hat{x}_{i+1} - \hat{x}_i|}{|\hat{x}_i|} \times 100. \quad (4)$$

That is, the iterative process is determined to have sufficiently converged once the percent change between the reconstructions, $\hat{x}$, from two subsequent iterations is less than 1%. Posing tighter restrictions than a percent change of less than 1% on the definition of convergence did not visually create a significant increase in image quality.

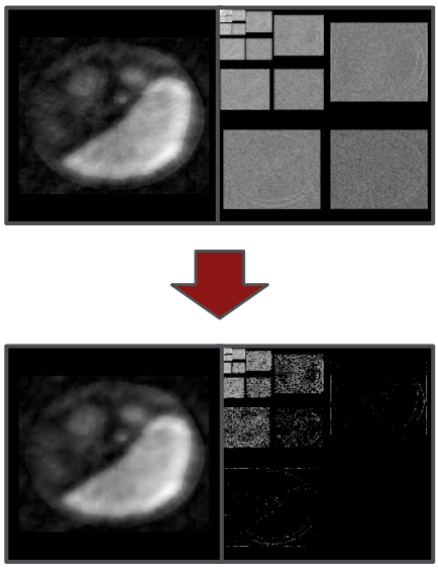

Fig. 4. Thresholding low energy wavelet coefficients results in a sparse approximation to the original data.

This iterative process is applied separately to each individual coil at every time step and repeated ten times, using threshold values linearly decreasing from 40% of the largest wavelet coefficient to 25% of the largest wavelet coefficient. By gradually increasing the amount of coefficients that are removed from the reconstruction fewer overall iterations are needed in order to achieve convergence [4].

Figure 5 shows the qualitative results after performing just this iterative reconstruction on the full (not eigencoil) parallel DCE phantom data. Comparison between the reconstruction without compressed sensing (figure 5 left) and the reconstruction after the iterative wavelet thresholding

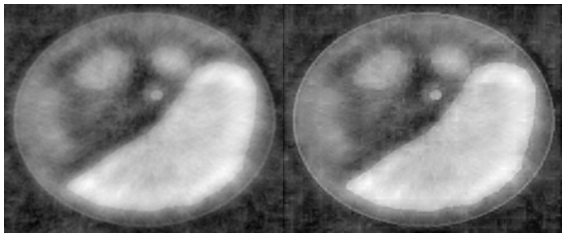

Fig. 5. Reconstruction without compressed sensing (left) and reconstruction after wavelet-based POCS type algorithm.

(figure 5 right) shows that the edges of the structures are much sharper, and the outlines of the most-heavily noise-masked structures (left and top of image) can be fully discerned.

2) Temporal Compressed Sensing: While the previously-described spatial compressed sensing method can be applied with good results to nearly any MR dataset, DCE data has an additional significant sparse representation in time. The only change between a DCE MR image from one time step and the corresponding cross-section at the subsequent time step will be a change is the intensity of some structures as the contrast agent flows through the body. As the majority of the image features are static throughout time, the time series data contains a large amount of redundant information describing the anatomic structures.

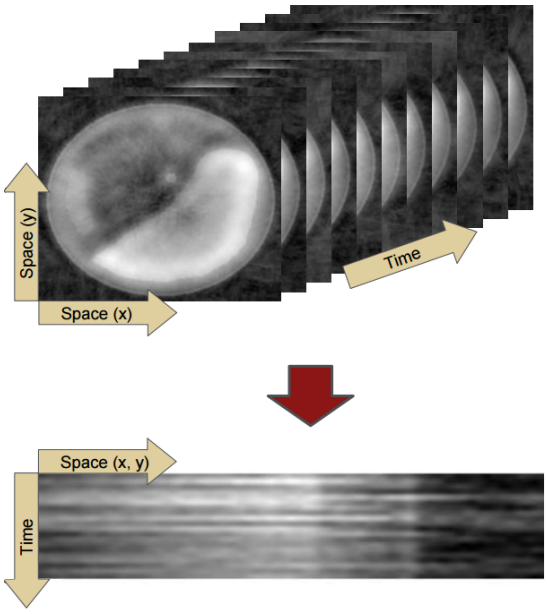

Fig. 6. Three-dimensional data (above) is compressed to the two-dimensional Casorati matrix with the original spatial dimensions x and y merged to one dimension and with time, t, as the second dimension.

A procedure similar to that described in section II.c1 is implemented in order to utilize a spar-

sifying transform to exploit this temporal data redundancy while simultaneously using SPIRiT to perform additional spatial compressed sensing.

Before applying the sparsifying transform, the data from each coil is first arranged into a two-dimensional Casorati matrix, as shown in figure 6. The original three-dimensional coil data (two spatial dimensions and one temporal dimension) is arranged such that resulting matrix has a column for every spatial coordinate with the temporal data for each spatial coordinate contained in the column.

Multiple sparsifying transformations were applied to the Casorati representation, including the wavelet transform used in II.c1 as well as a thresholding of Fourier transform coefficients and a low-rank approximation, as suggested in [1]. The low-rank approximation was found to be significantly more effective than the either evaluated methods. The Casorati matrix has very low effective rank, and an accurate sparse reconstruction can be performed using a small number of significant singular values [1]. The algorithm for doing so is provided here.

1) Apply SPIRiT
2) In image space, reshape data into Casorati matrix representation
3) Calculate the singular value decomposition of the Casorati matrix
4) Threshold out non-significant singular values
5) Construct a low-rank approximation to the original data using the significant singular values
6) Repeat 1 through 5 until sufficient convergence is achieved

My implementation of this algorithm defines sufficient convergence as either the percent change between two subsequent iterations effectively reaching zero ($< 10^{-9}$) or the percent change fluctuating between increasing and decreasing, as this indicates that the procedure may not be able to find a more satisfactory convergence. The results of applying SPIRiT with the low-rank constraint to the phantom DCE data is shown in figure 7 (right).

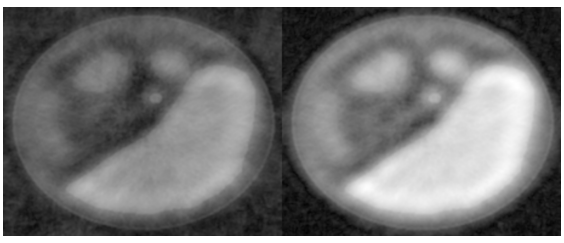

Fig. 7. Reconstruction without compressed sensing (left) and reconstruction after low-rank-approximation-based POCS type algorithm.

III. Results

The method presented here is capable of reconstructing original DCE MRI data from aggressively undersampled data with good accuracy. Below, I give the results of my implementation on three DCE datasets, with side-by-side comparisons to the result obtained from a basic inverse Fourier transform reconstruction. As stated previously, all multi-coil reconstructions presented in this report were formed using equation 1.

In each test, significant denoising has been achieved, and the resolution of the reconstruction is high enough that all major and most minor anatomic structures can be discerned. Figures 8, 9, and 10 show the reconstructions only at the final time point in the time series. Figure 11 gives the full twenty reconstructed time-series images for the dataset shown in figure 9.

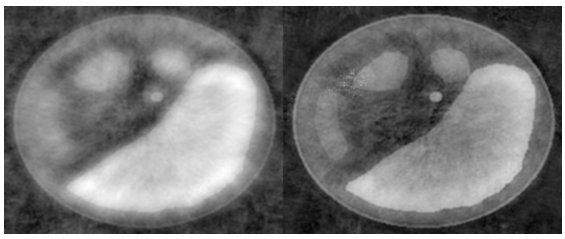

Fig. 8. $16\times$ variable-density Poisson disk undersampled DCE phantom reconstruction. All six primary structures can be made out with strong edges, and significant denoising has been achieved.

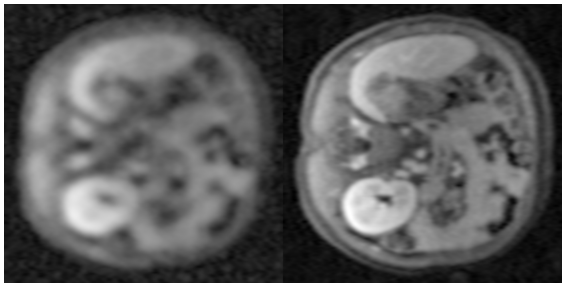

Fig. 9. $18\times$ variable-density Poisson disk undersampled pediatric abdominal cross-section DCE MRI. Note the definition of smaller features and the sharpness of feature edges.

IV. Related Work

This project combines the algorithms for spatio-temporal compressed sensing laid out by Michael Lustig [3] [4] [5] [6] with the procedure for DCE MRI applications outlined by Tao Zhang [1]. My proposed method implements the algorithm described in [1] with the addition of the wavelet thresholding step detailed in section II.c1.

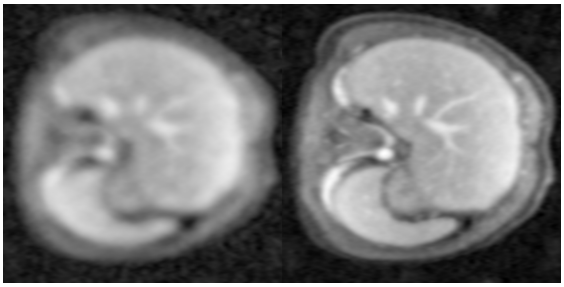

Fig. 10. $18\times$ variable-density Poisson disk undersampled pediatric abdominal cross-section DCE MRI. Note the definition of the vasculature in the liver (largest structure, upper-right).

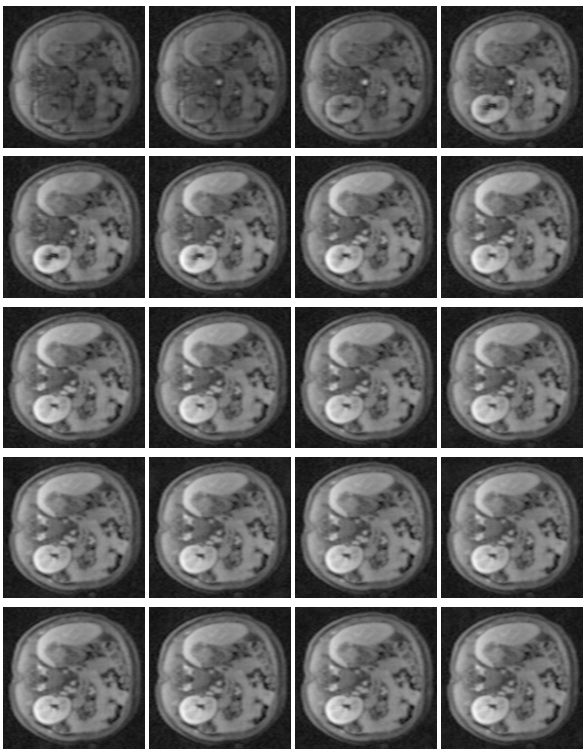

Fig. 11. Full twenty-frame reconstruction of the $18\times$ variable-density Poisson disk undersampled pediatric abdominal cross-section shown in figure 9. Time increases from left-to-right and top-to-bottom. Note the intensity changes of the anatomic structures, which show the tissue uptake of the contrast agent.

V. Future Work

A. Compressed Sensing: Wavelets

Computing numerous forward and inverse wavelet transforms is computationally expensive. As such, only a relatively small number of iterations of the POCS type wavelet solution were run in my implementation. With more extensive computation power, it should be possible to converge to a more accurate iterative solution by iterating the procedure many more times over a finer variation of coefficient thresholds.

B. Parameter Choice

A more rigorous evaluation of the Poisson disk undersampling scheme implemented should be performed. A more carefully selected selection of k-space points with a precisely controlled variation over time, as implemented in [1] should be performed with an evaluation of the performance of varying the distance between points and the overall density variation of sample points as a function of spatial frequency.

REFERENCES

[1] T. Zhang, J. Y. Cheng, A. G. Potnick, R. A. Barth, M. T. Alley, M. Uecker, M. Lustig, J. M. Pauly, and S. S. Vasanawala, "Fast pediatric 3d free-breathing abdominal dynamic contrast enhanced mri with high spatiotemporal resolution," *Journal of Magnetic Resonance Imaging*, 2015.

[2] M. Lustig and J. M. Pauly, "SPIRiT, iterative self-consistent parallel imaging reconstruction from arbitrary *k*-space," *Magnetic Resonance in Medicine*, vol. 64, pp. 457–471, 2010.

[3] M. Lustig, J. M. Santos, D. Donoho, and J. M. Pauly, "k-t SPARSE: High frame rate dynamic MRI exploiting spatio-temporal sparsity," in *Proceedings of the 14th Annual Meeting of ISMRM, Seattle*, p. 2420, 2006.

[4] M. Lustig, "EE369C: Assignment 6," Accessed: December 1, 2015. http://web.stanford.edu/class/ee369c/Assignments/assignment_6.pdf.

[5] M. Lustig, D. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, pp. 1182–1195, 2007.

[6] M. Lustig, J. H. Lee, D. L. Donoho, and J. M. Pauly, "Faster imaging with randomly perturbed, undersampled spirals and l_1 reconstructions," in *Proceedings of the 12th Annual Meeting of ISMRM, Miami Beach*, p. 685, 2005.

[7] J. Pauly, "EE369C: Medical Image Reconstruction Lecture Notes," Accessed: December 1, 2015. http://web.stanford.edu/class/ee369c/Notes.html.