Sparse MRI: The Application of Compressed Sensing for Rapid MR Imaging

Michael Lustig,[1]* David Donoho,[2] and John M. Pauly[1]

The sparsity which is implicit in MR images is exploited to significantly undersample *k*-space. Some MR images such as angiograms are already sparse in the pixel representation; other, more complicated images have a sparse representation in some transform domain–for example, in terms of spatial finite-differences or their wavelet coefficients. According to the recently developed mathematical theory of compressed-sensing, images with a sparse representation can be recovered from randomly undersampled *k*-space data, provided an appropriate nonlinear recovery scheme is used. Intuitively, artifacts due to random undersampling add as noise-like interference. In the sparse transform domain the significant coefficients stand out above the interference. A nonlinear thresholding scheme can recover the sparse coefficients, effectively recovering the image itself. In this article, practical incoherent undersampling schemes are developed and analyzed by means of their aliasing interference. Incoherence is introduced by pseudo-random variable-density undersampling of phase-encodes. The reconstruction is performed by minimizing the ℓ^1 norm of a transformed image, subject to data fidelity constraints. Examples demonstrate improved spatial resolution and accelerated acquisition for multislice fast spin-echo brain imaging and 3D contrast enhanced angiography. Magn Reson Med 58:1182–1195, 2007. © 2007 Wiley-Liss, Inc.

Key words: compressed sensing; compressive sampling; random sampling; rapid MRI; sparsity; sparse reconstruction; nonlinear reconstruction

INTRODUCTION

Imaging speed is important in many MRI applications. However, the speed at which data can be collected in MRI is fundamentally limited by physical (gradient amplitude and slew-rate) and physiological (nerve stimulation) constraints. Therefore, many researches are seeking for methods to reduce the amount of acquired data without degrading the image quality.

When *k*-space is undersampled, the Nyquist criterion is violated, and Fourier reconstructions exhibit aliasing artifacts. Many previous proposals for reduced data imaging try to mitigate undersampling artifacts. They fall in three groups: (a) Methods generating artifacts that are incoherent or less visually apparent, at the expense of reduced apparent SNR (1–5); (b) Methods exploiting redundancy in *k*-space, such as partial-Fourier, parallel imaging, etc. (6–8); (c) Methods exploiting either spatial or temporal redundancy or both (9–13).

In this article we aim to exploit the sparsity which is implicit in MR images, and develop an approach combining elements of approaches *a* and *c*. By implicit sparsity we mean transform sparsity, i.e., the underlying object we aim to recover happens to have a sparse representation in a known and fixed mathematical transform domain. To begin with, consider the identity transform, so that the transform domain is simply the image domain itself. Here sparsity means that there are relatively few significant pixels with nonzero values. For example, angiograms are extremely sparse in the pixel representation. More complex medical images may not be sparse in the pixel representation, but they do exhibit transform sparsity, since they have a sparse representation in terms of spatial finite differences, in terms of their wavelet coefficients, or in terms of other transforms.

Sparsity is a powerful constraint, generalizing the notion of finite object support. It is well understood why support constraints in image space (i.e., small FOV or band-pass sampling) enable sparser sampling of *k*-space. Sparsity constraints are more general because nonzero coefficients do not have to be bunched together in a specified region. Transform sparsity is even more general because the sparsity needs only to be evident in some transform domain, rather than in the original image (pixel) domain. Sparsity constraints, under the right circumstances, can enable sparser sampling of *k*-space as well (14,15).

The possibility of exploiting transform sparsity is motivated by the widespread success of data compression in imaging. Natural images have a well-documented susceptibility to compression with little or no visual loss of information. Medical images are also compressible, though this topic has been less thoroughly studied. Underlying the most well-known image compression tools such as JPEG, and JPEG-2000 (16) are the discrete cosine transform (DCT) and wavelet transform. These transforms are useful for image compression because they transform image content into a vector of sparse coefficients; a standard compression strategy is to encode the few significant coefficients and store them, for later decoding and reconstruction of the image.

The widespread success of compression algorithms with real images raises the following questions: Since the images we intend to acquire will be compressible, with most transform coefficients negligible or unimportant, is it really necessary to acquire all that data in the first place? Can we not simply measure the compressed information directly

[1]Magnetic Resonance Systems Research Laboratory, Department of Electrical Engineering, Stanford University, Stanford, California

[2]Statistics Department, Stanford University, Stanford, California

Contract grant sponsor: NIH; Contract grant numbers: R01 HL074332, R01 HL 067161, R01 HL075803; Contract grant sponsor: NSF; Contract grant number: DMS 0505303; Contract grant sponsor: GE Healthcare.

*Correspondence to: Michael Lustig, Room 210, Packard Electrical Engineering Bldg, Stanford University, Stanford, CA 94305-9510. E-mail: mlustig@mrsrl.stanford.edu

Received 22 December 2006; revised 18 July 2007; accepted 20 July 2007.

DOI 10.1002/mrm.21391

Published online 29 October 2007 in Wiley InterScience (www.interscience.wiley.com).

© 2007 Wiley-Liss, Inc.

from a small number of measurements, and still reconstruct the same image which would arise from the fully sampled set? Furthermore, since MRI measures Fourier coefficients, and not pixels, wavelet, or DCT coefficients, the question is whether it is possible to do the above by measuring only a subset of *k*-space.

A substantial body of mathematical theory has recently been published establishing the possibility to do exactly this. The formal results can be found by searching for the phrases *compressed sensing* (CS) or *compressive sampling* (14,15). According to these mathematical results, if the underlying image exhibits transform sparsity, and if *k*-space undersampling results in incoherent artifacts in that transform domain, then the image can be recovered from randomly undersampled frequency domain data, provided an appropriate nonlinear recovery scheme is used.

In this article we aim to develop a framework for using CS in MRI. To keep the discussion as short and simple as possible, we focus this work only on Cartesian sampling. Since most product pulse sequences in the clinic today are Cartesian, the impact of Cartesian CS can be substantial. We keep in mind though, that non-Cartesian CS has great potential and may be even more advantageous than Cartesian for some applications. Some very promising results for radial and spiral imaging have been presented by (17–21).

THEORY

Compressed Sensing

CS was first proposed in the literature of Information Theory and Approximation Theory in an abstract general setting. One measures a small number of random linear combinations of the signal values–much smaller than the number of signal samples nominally defining it. The signal is reconstructed with good accuracy from these measurements by a nonlinear procedure. In MRI we look at a special case of CS, where the sampled linear combinations are simply individual Fourier coefficients (*k*-space samples). In that setting, CS is claimed to be able to make accurate reconstructions from a small subset of *k*-space rather than an entire *k*-space grid.

The CS approach requires that: (a) the desired image have a sparse representation in a known transform domain (i.e., is compressible), (b) the aliasing artifacts due to *k*-space undersampling be incoherent (noise like) in that transform domain. (c) a nonlinear reconstruction be used to enforce both sparsity of the image representation and consistency with the acquired data. To help keep in mind these ingredients, consider Fig. 1, which depicts relationships among some of these main concepts. It shows the image, the *k*-space and the transform domains, and the operators connecting these domains and the requirements for CS.

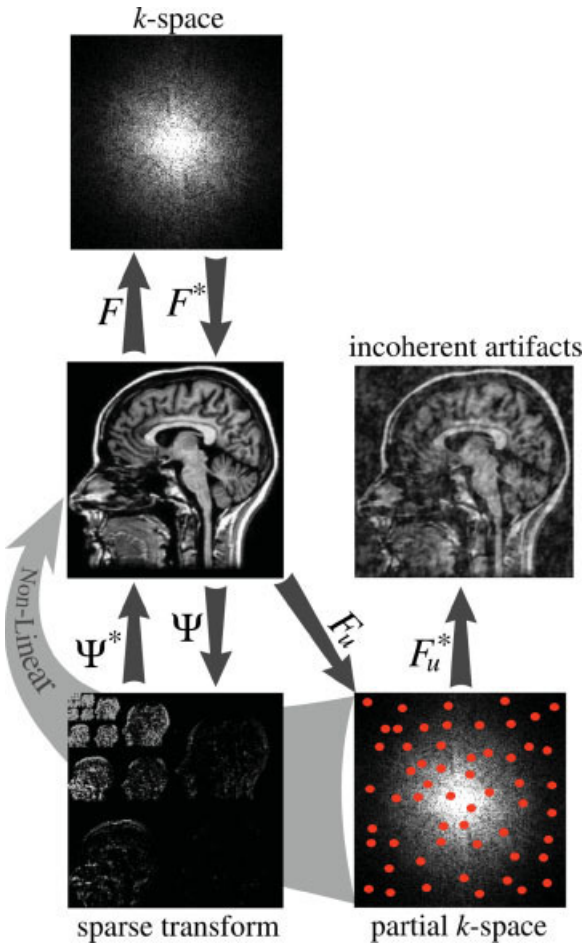

FIG. 1. Illustration of the domains and operators used in the paper as well as the requirements of CS: sparsity in the transform domain, incoherence of the undersampling artifacts, and the need for nonlinear reconstruction that enforces sparsity. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

A Simple, Intuitive Example of Compressed Sensing

To get intuition for the importance of incoherence and the feasibility of CS in MRI, consider the example in Fig. 2. A sparse 1D signal (Fig. 2a), 256 samples long, is undersampled in *k*-space (Fig. 2b) by a factor of eight. Here, the sparse transform is simply the identity. Later, we will consider the case where the transform is nontrivial.

Equispaced *k*-space undersampling and reconstruction by zero-filling results in coherent aliasing, a superposition of shifted replicas of the signal as illustrated in Fig. 2c. In this case, there is an inherent ambiguity; it is not possible to distinguish between the original signal and its replicas, as they are all equally likely.

Random undersampling results in a very different situation. The zero-filling Fourier reconstruction exhibits incoherent artifacts that actually behave much like additive random noise (Fig. 2d). Despite appearances, the artifacts are not noise; rather, undersampling causes leakage of energy away from each individual nonzero coefficient of the original signal. This energy appears in other reconstructed signal coefficients, including those which had been zero in the original signal.

It is possible, if all the underlying original signal coefficients are known, to calculate this leakage analytically.

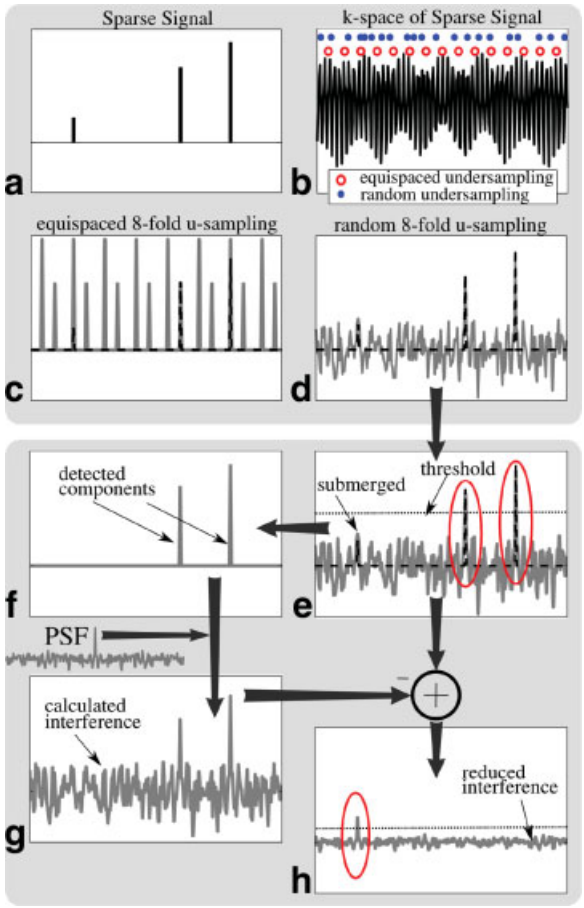

FIG. 2. An intuitive reconstruction of a sparse signal from pseudo-random *k*-space undersampling. A sparse signal (a) is 8-fold undersampled in *k*-space (b). Equispaced undersampling results in coherent signal aliasing (c) that cannot be recovered. Pseudo-random undersampling results in incoherent aliasing (c). Strong signal components stick above the interference, are detected (e) and recovered (f) by thresholding. The interference of these components is computed (g) and subtracted (h), lowering the total interference level and enabling recovery of weaker components. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

This observation enables the signal in Fig. 2d to be accurately recovered although it was 8-fold undersampled. An intuitive plausible recovery procedure is illustrated in Fig. 2e–h. It is based on thresholding, recovering the strong components, and calculating the interference caused by them and subtracting it. Subtracting the interference of the strong components reduces the total interference level and enables recovery of weaker, previously submerged components. By iteratively repeating this procedure, one can recover the rest of the signal components. A recovery procedure along these lines was proposed by Donoho et al. (Sparse Solution of Underdetermined Linear Equations by Stagewise Orthogonal Matching Pursuit, 2006, Stanford University, Statistics Department, technical report #2006-02) as a fast approximate algorithm for CS reconstruction. A similar approach of recovery of MR images was proposed in Ref. (22).

Sparsity

Sparsifying Transform

A sparsifying transform is an operator mapping a vector of image data to a sparse vector. In recent years, there has been extensive research in sparse image representation. As a result, we currently possess a library of diverse transformations that can sparsify many different type of images (23).

For example, piecewise constant images can be sparsely represented by spatial finite-differences (i.e, computing the differences between neighboring pixels); indeed, away from boundaries, the differences vanish. Real-life MR images are of course not piecewise smooth. But in some problems, where boundaries are the most important information (angiograms for example) computing finite-differences results in a sparse representation.

Natural, real-life images are known to be sparse in the discrete cosine transform (DCT) and wavelet transform domains (16). The DCT is central to the JPEG image compression standard and MPEG video compression, and is used billions of times daily to represent images and videos. The wavelet transform is used in the JPEG-2000 image compression standard (16). The wavelet transform is a multiscale representation of the image. Coarse-scale wavelet coefficients represent the low resolution image components and fine-scale wavelet coefficients represent high resolution components. Each wavelet coefficient carries both spatial position and spatial frequency information at the same time (see top Fig. 4b for a spatial position and spatial frequency illustrations of a mid-scale wavelet coefficient).

Since computing finite-differences of images is a high-pass filtering operation, the finite-differences transform can also be considered as computing some sort of fine-scale wavelet transform (without computing coarser scales).

Sparsity is not limited only to the spatial domain. Dynamic images are extremely sparse in the temporal dimension. Dynamic sparsity is beyond our scope; some preliminary results of dynamic CS imaging are reported in Refs. (24) and (25).

The Sparsity of MR Images

The transform sparsity of MR images can be demonstrated by applying a sparsifying transform to a fully sampled image and reconstructing an approximation to the image from a subset of the largest transform coefficients. The sparsity of the image is the percentage of transform coefficients sufficient for diagnostic-quality reconstruction. Of course the term "diagnostic quality" is subjective. Nevertheless, for specific applications, it is possible to get an empirical sparsity estimate by performing a clinical trial and evaluating reconstructions of many images quantitatively or qualitatively.

To illustrate this, we performed such an experiment on two representative MR images: an angiogram of a leg and a brain image. The images were transformed by each transform of interest and reconstructed from several subsets of the largest transform coefficients. The results are depicted in Fig. 3. The left column images show the magnitude of the transform coefficients; they illustrate that indeed the transform coefficients are sparser than the images

themselves. The DCT and the wavelet transforms have similarly good performance with a slight advantage for the wavelet transform for both brain and angiogram images at reconstructions involving 5–10% of the coefficients. The finite-difference transform does not sparsify the brain image well. Nevertheless, finite differences do sparsify angiograms because they primarily detect the boundaries of the blood vessels, which occupy less than 5% of the spatial domain.

Incoherent Sampling: "Randomness is too Important to be Left to Chance[1]"

Incoherent aliasing interference in the sparse transform domain is an essential ingredient for CS. This can be well understood from our previous simple 1D example. In the original CS papers (14,15), sampling a completely random subset of k-space was chosen to simplify the mathematical proofs and in particular to guarantee a very high degree of incoherence.

Random point k-space sampling in all dimensions is generally impractical as the k-space trajectories have to be relatively smooth because of hardware and physiological considerations. Instead, we aim to design a practical incoherent sampling scheme that mimics the interference properties of pure random undersampling as closely as possible yet allows rapid collection of data.

There are numerous ways to design incoherent sampling trajectories. To focus and simplify the discussion, in this article, we consider only the case of Cartesian grid sampling where the sampling is restricted to undersampling the phase-encodes and fully sampled readouts. Alternative sampling trajectories are possible and some very promising results have been presented by Refs. (19–21) (radial imaging), and by Refs. (17) and (18) (spiral imaging).

We focus on Cartesian sampling because it is by far the most widely used in practice. It is simple and also highly robust to numerous sources of imperfection. Nonuniform undersampling of phase encodes in Cartesian imaging has been proposed in the past as an acceleration method because it produces incoherent artifacts (1,3,5)–exactly what we are looking for. Undersampling phase-encode lines offers pure randomness in the phase-encode dimensions, and a scan time reduction that is exactly proportional to the undersampling. Finally, implementation of such an undersampling scheme is simple and requires only minor modifications to existing pulse sequences.

Point Spread Function and Transform Point Spread Function Analysis

The point spread function (PSF) is a natural tool to measure incoherence. Let $\mathcal{F}_u$ be the undersampled Fourier operator and let e_i be the ith vector of the natural basis (i.e, having "1" at the ith location and zeroes elsewhere). Then $\mathrm{PSF}(i;j) = e_j^* \mathcal{F}_u^* \mathcal{F}_u e_i$ measures the contribution of a unit-intensity pixel at the ith position to a pixel at the jth position. Under Nyquist sampling there is no interference between pixels and $\mathrm{PSF}(i;j)|_{i\neq j} = 0$. Undersampling causes pixels to interfere and $\mathrm{PSF}(i;j)|_{i\neq j}$ to assume nonzero values. A simple measure to evaluate the incoherence is the maximum of the sidelobe-to-peak ratio (SPR), $\max_{i\neq j} \left| \frac{\mathrm{PSF}(i,j)}{\mathrm{PSF}(i,i)} \right|$.

The PSF of pure 2D random sampling, where samples are chosen at random from a Cartesian grid, offers a standard for comparison. In this case $\mathrm{PSF}(i;j)|_{i\neq j}$ looks random as illustrated in Fig. 4a. Empirically, the real and the imaginary parts separately behave much like zero-mean random white Gaussian noise. The standard deviation of the observed SPR depends on the number, N, of samples taken and the number, D, of grid points defining the underlying image. For a constant sampling reduction factor $p = \frac{D}{N}$ the standard deviation obeys the formula:

$$\sigma_{\mathrm{SPR}} = \sqrt{\frac{p-1}{D}}. \qquad [1]$$

A derivation of Eq. [1] is given in Appendix II.

The MR images of interest are typically sparse in a transform domain rather than the usual image domain. In such a setting, incoherence is analyzed by generalizing the notion of PSF to Transform Point Spread Function (TPSF) which measures how a single transform coefficient of the underlying object ends up influencing other transform coefficients of the measured undersampled object.

Let Ψ be an orthogonal sparsifying transform (nonorthogonal TPSF analysis is beyond our scope and is not discussed here). The $\mathrm{TPSF}(i;j)$ is given by the following equation,

$$\mathrm{TPSF}(i;j) = e_j^* \Psi \mathcal{F}_u^* \mathcal{F}_u \Psi^* e_i. \qquad [2]$$

In words, a single point in the transform space at the ith location is transformed to the image space and then to the Fourier space. The Fourier space is subjected to undersampling, then transformed back to the image space. Finally, a return is made to the transform domain and the jth location of the result is selected. An example using an orthogonal wavelet transform is illustrated by Fig. 4b. The size of the sidelobes in $\mathrm{TPSF}(i;j)|_{i\neq j}$ are used to measure the incoherence of a sampling trajectory. We would like $\mathrm{TPSF}(i;j)|_{i\neq j}$ to be as small as possible, and have random noise-like statistics.

Single-slice 2DFT, multislice 2DFT, and 3DFT Imaging

Equipped with the PSF and TPSF analysis tools, we consider three cases of Cartesian sampling: 2DFT, multislice 2DFT, and 3DFT. In single-slice 2DFT, only the phase encodes are undersampled and the interference spreads only along a single dimension. The interference standard deviation as calculated in Eq. [1] is $D^{1/4}$ times larger than the theoretical pure random 2D case for the same acceleration–(16 times for a 256×256 image). Therefore, in 2DFT one can expect relatively modest accelerations because mostly 1D sparsity is exploited.

In multislice 2DFT we sample in a hybrid k-space versus image space ($k_y - z$ space). Undersampling differently the phase-encodes of each slice randomly undersamples the $k_y - z$ space. This can reduce the peak sidelobe in the TPSF of some appropriate transforms, such as wavelets, as

[1]Robert R. Coveyou, Oak Ridge National Laboratory.

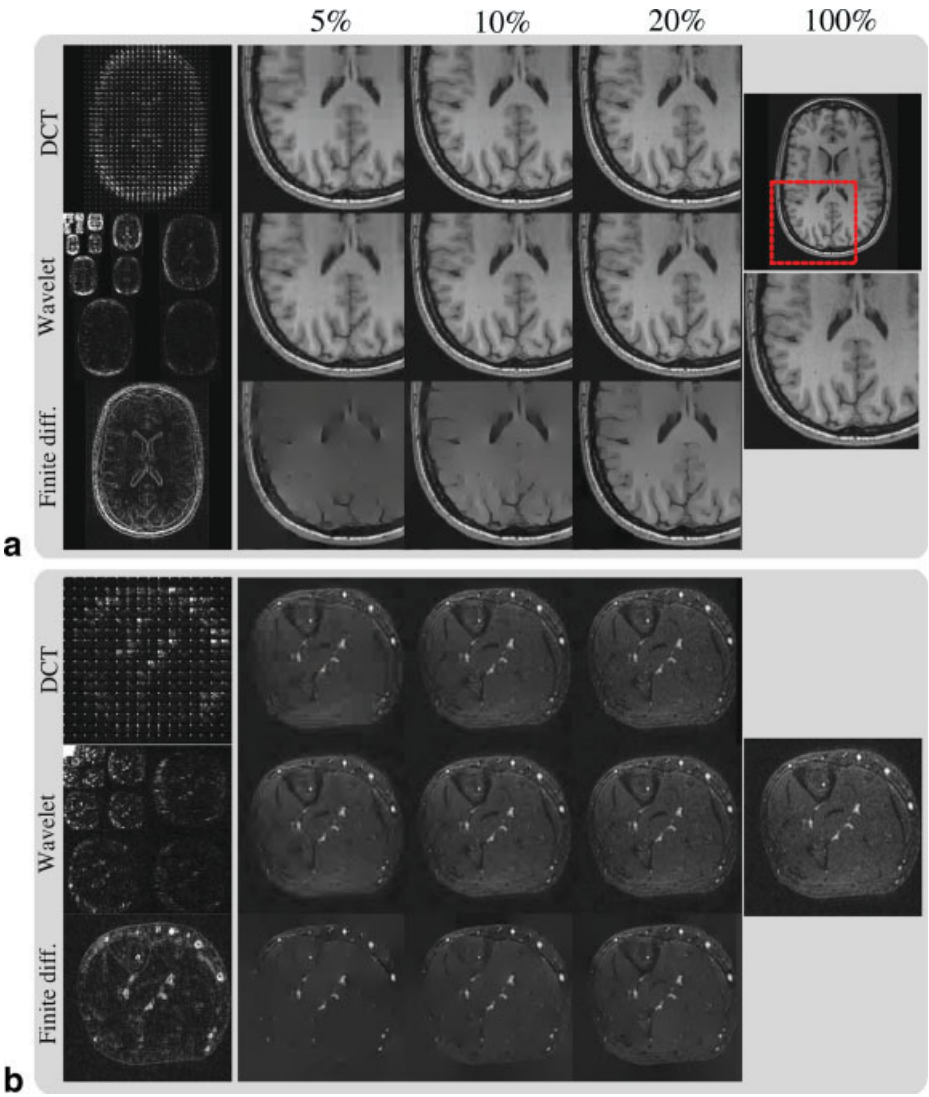

FIG. 3. Transform-domain sparsity of images. (a) Axial T_1 weighted brain image; (b) axial 3D contrast enhanced angiogram of the peripheral leg. The DCT, wavelet, and finite-differences transforms were calculated for all the images (Left column). The images were then reconstructed from a subset of 5, 10, and 20% of the largest transform coefficients.

long as the transform is also applied in the slice dimension. Hence, it is possible to exploit some of the sparsity in the slice dimension as well. Figure 5a, b shows that undersampling each slice differently has reduced peak sidelobes in the TPSF compared to undersampling the slices the same way. However, it is important to mention that for wavelets, randomly undersampling in the hybrid $k_y - z$ space is not as effective, in terms of reducing the peak sidelobes, as randomly undersampling in a pure 2D k-space (Fig. 5c). The method of multislice 2DFT will work particularly well

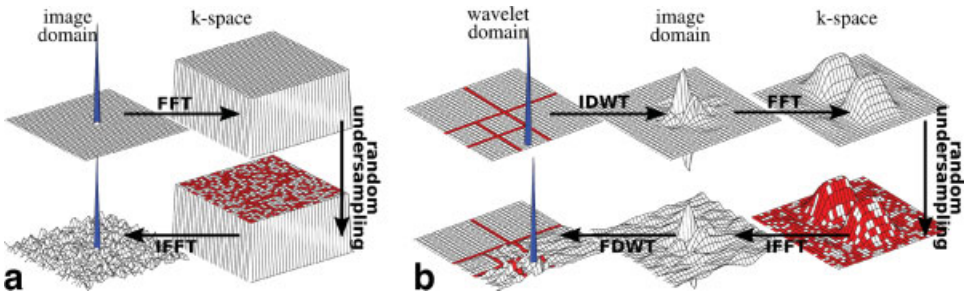

FIG. 4. (a) The PSF of random 2D k-space undersampling. (b) The wavelet TPSF of random 2D Fourier undersampling. FDWT and IDWT stand for forward and inverse discrete wavelet transform. Wavelet coefficients are band-pass filters and have limited support both in space and frequency. Random k-space undersampling results in incoherent interference in the wavelet domain. The interference spreads mostly within the wavelet coefficients of the same scale and orientation.

when the slices are thin and finely spaced. When the slices are thick and with gaps, there is little spatial redundancy in the slice direction and the performance of the reconstruction would be reduced to the single-slice 2DFT case. Undersampling with CS can be used to bridge gaps or acquire more thinner slices without compromising the scan time.

Randomly undersampling the 3DFT trajectory is our preferred method. Here, it is possible to randomly undersample the 2D phase encode plane ($k_y - k_z$) and achieve the theoretical high degree of 2D incoherence. Additionally, 2D sparsity is fully exploited, and images have a sparser representation in 2D. Three-dimensional imaging is particularly attractive because it is often time consuming and scan time reduction is a higher priority than 2D imaging. Figure 5c illustrates the proposed undersampled 3DFT trajectory and its wavelet TPSF. The peak interference of the wavelet coefficients is significantly reduced compared to multislice and plain 2DFT undersampling.

Variable Density Random Undersampling

Our incoherence analysis so far assumes the few nonzeros are scattered at random among the entries of the transform domain representation. Representations of natural images exhibit a variety of significant nonrandom structures. First, most of the energy of images is concentrated close to the k-space origin. Furthermore, using wavelet analysis one can observe that coarse-scale image components tend to be less sparse than fine-scale components. This can be seen in the wavelet decomposition of the brain and angiogram images of Fig. 3, left column.

These observations show that, for a better performance with "real images," one should be undersampling less near the k-space origin and more in the periphery of k-space. For example, one may choose samples randomly with sampling density scaling according to a power of distance from the origin. Empirically, using density powers of 1–6 greatly reduces the total interference and, as a result, iterative algorithms converge faster with better reconstruction. The optimal sampling density is beyond the scope of this article, and should be investigated in future research.

Variable-density sampling schemes for Cartesian, radial (radial has natural linear density), and spiral imaging have been proposed in the past by (1–5) because the aliasing appears incoherent. In such schemes, high energy low-frequency image components alias less than lower energy higher-frequency components, and the interference appears as white noise in the image domain. This is exactly the desired case in CS, only in CS it is also possible to remove the aliasing interference without degrading the image quality.

How Many Samples to Acquire?

A theoretical bound on the number of Fourier sample points that need be collected with respect to the number of sparse coefficients is derived in Refs. (14) and (15). However, we as well as other researchers have observed that in practice, for a good reconstruction, the number of k-space samples should be roughly two to five times the number of sparse coefficients (The number of sparse coefficients can be calculated in the same way as in the The Sparsity of MR Images section). Our results, presented in this article, support this claim. Similar observations were reported by Candès et al. (26) and by Tsaig and Donoho (27).

Monte-Carlo Incoherent Sampling Design

Finding an optimal sampling scheme that maximizes the incoherence for a given number of samples is a combinatorial optimization problem and might be considered intractable. However, choosing samples at random often results in a good, incoherent, near-optimal solution. Therefore, we propose the following Monte-Carlo design procedure: choose a grid size based on the desired resolution and FOV of the object. Undersample the grid by constructing a probability density function (pdf) and randomly draw indices from that density. Variable density sampling of k-space is controlled by the pdf construction. A plausible choice is diminishing density according to a power of distance from the origin as previously discussed. Because the procedure is random, one might accidentally choose a sampling pattern with a "bad" TPSF. To prevent such situation, repeat the procedure many times, each time measure the peak interference in the TPSF of the resulting sampling pattern. Finally, choose the pattern with the lowest peak interference. Once a sampling pattern is determined it can be used again for future scans.

Image Reconstruction

We now describe in more detail the processes of nonlinear image reconstruction appropriate to the CS setting. Suppose the image of interest is a vector m, let Ψ denote the linear operator that transforms from pixel representation into a sparse representation, and let $\mathcal{F}_u$ be the undersampled Fourier transform, corresponding to one of the k-space undersampling schemes discussed earlier. The reconstruction is obtained by solving the following constrained optimization problem:

$$\begin{aligned} &minimize \quad \|\Psi m\|_1 \\ &s.t. \quad \|\mathcal{F}_u m - y\|_2 < \epsilon \end{aligned} \qquad [3]$$

Here m is the reconstructed image, where y is the measured k-space data from the scanner, and ϵ controls the fidelity of the reconstruction to the measured data. The threshold parameter ϵ is usually set below the expected noise level.

The objective function in Eq. [3] is the ℓ_1 norm, which is defined as $\|x\|_1 = \sum_i |x_i|$. Minimizing $\|\Psi m\|_1$ promotes sparsity (28). The constraint $\|\mathcal{F}_u m - y\|_2 < \epsilon$ enforces data consistency. In words, among all solutions which are consistent with the acquired data, Eq. [3] finds a solution which is compressible by the transform Ψ.

When finite-differences is used as a sparsifying transform, the objective in Eq. [3] is often referred to as Total-Variation (TV) (29), since it is the sum of the absolute variations in the image. The objective then is usually written as TV(m). Even when using other sparsifying transforms in the objective, it is often useful to include a TV penalty as well (27). This can be considered as requiring the image to be sparse by both the specific transform and

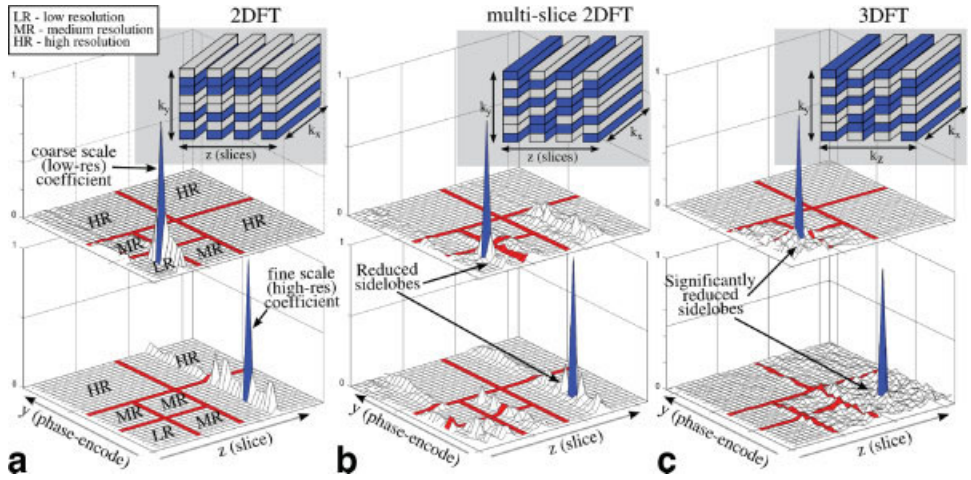

FIG. 5. Transform point spread function (TPSF) analysis in the wavelet domain. The k-space sampling patterns and the associated TPSF of coarse-scale and fine-scale wavelet coefficients are shown. (a) Random phase encode undersampling spreads the interference only in 1D and mostly within the same wavelet scale. The result is relatively high peak interference. (b) Sampling differently for each slice, i.e., randomly undersampling the $k_y - z$ plane causes the interference to spread to nearby slices and to other wavelets scales and reduces its peak value. (c) Undersampling the phase encode plane, i.e., $k_y - k_z$ spreads the interference in 2D and results in the lowest peak interference. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

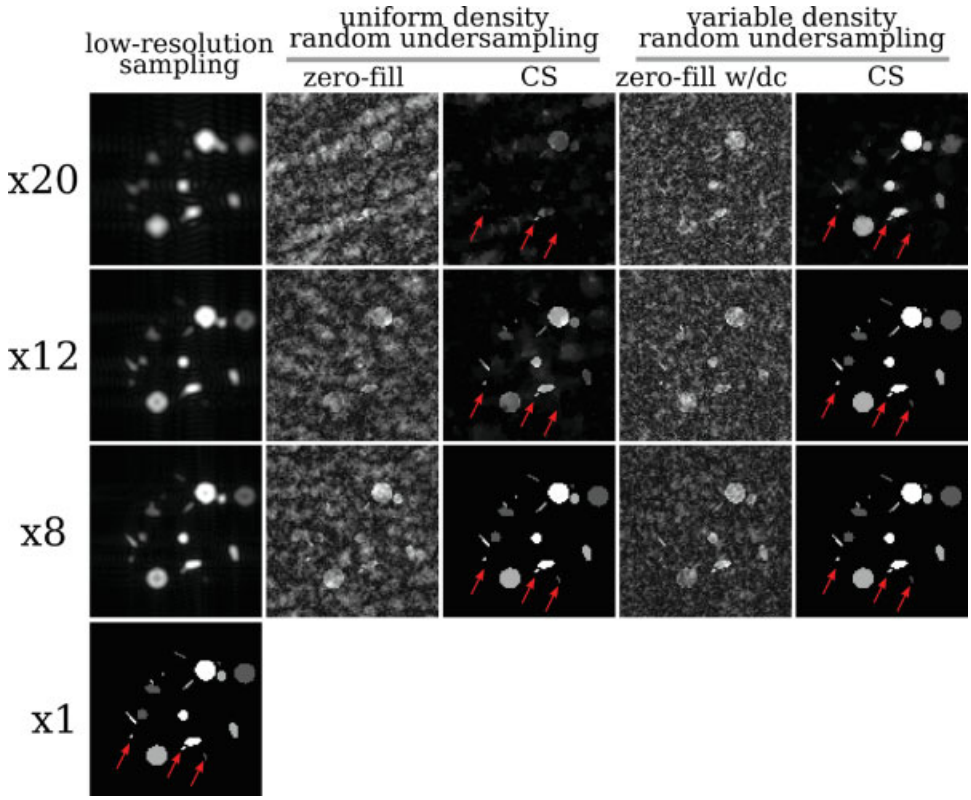

FIG. 6. Simulation: Reconstruction artifacts as a function of acceleration. The LR reconstructions exhibit diffused boundaries and loss of small features. The ZF-w/dc reconstructions exhibit an significant increase of apparent noise due to incoherent aliasing, the apparent noise appears more "white" with variable density sampling. The CS reconstructions exhibit perfect reconstruction at 8- and 12-fold (only var. dens.) accelerations. With increased acceleration there is loss of low-contrast features and not the usual loss of resolution. The reconstructions from variable density random undersampling significantly outperforms the reconstructions from uniform density random undersampling. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

finite-differences at the same time. In this case Eq. [3] is written as

$$\begin{aligned} \textit{minimize} \quad & \|\Psi m\|_1 + \alpha TV(m) \\ s.t. \quad & \|\mathcal{F}_u m - y\|_2 < \epsilon, \end{aligned}$$

where α trades Ψ sparsity with finite-differences sparsity.

The ℓ_1 norm in the objective is a crucial feature of the whole approach. Minimizing the ℓ_1 norm of an objective often results in a sparse solution. On the other hand, minimizing the ℓ_2 norm, which is defined as $\|x\|_2 = (\sum_i |x_i|^2)^{1/2}$ and commonly used for regularization because of its simplicity, does not result in a sparse solution and hence is not suitable for use as objective function in Eq. [3]. Intuitively, the ℓ_2 norm penalizes large coefficients heavily, therefore solutions tend to have many smaller coefficients–hence not be sparse. In the ℓ_1 norm, many small coefficients tend to carry a larger penalty than a few large coefficients, therefore small coefficients are suppressed and solutions are often sparse.

Special purpose methods for solving Eq. [3] have been a focus of research interest since CS was first introduced. Proposed methods include: interior point methods (28,30), projections onto convex sets (26), homotopy (Donoho et al., Fast solution of ℓ^1 minimization where the solution may be sparse 2006, Statistics Department, Stanford University, technical report #2006 18), iterative soft thresholding (31–33), and iteratively reweighted least squares (20,34). In the Appendix we describe our approach which is similar to (19,21,35), using nonlinear conjugate gradients and backtracking line-search.

It is important to mention that some of the above iterative algorithms for solving the optimization in Eq. [3] in effect perform thresholding and interference cancellation at each iteration. Therefore, there is a close connection between our previous simple intuitive example of interference cancellation and the more formal approaches that are described above.

Low-Order Phase Correction and Phase Constrained Partial *k*-space

In MRI, instrumental sources of phase errors can cause low-order phase variation. These carry no physical information, but create artificial variation in the image which makes it more difficult to sparsify, especially by finite differences. By estimating the phase variation, the reconstruction can be significantly improved. This phase estimate may be obtained using very low-resolution fully sampled *k*-space information. Alternatively, the phase is obtained by solving Eq. [3] to estimate the low-order phase, and repeating the reconstruction while correcting for the phase estimate.

The phase information is incorporated by a slight modification of Eq. [3],

$$\begin{aligned} \textit{minimize} \quad & \|\Psi m\|_1 \\ s.t. \quad & \|\mathcal{F}_u P m - y\|_2 < \epsilon \end{aligned} \qquad [4]$$

where P is a diagonal matrix whose entries give the estimated phase of each pixel.

METHODS

All experiments were performed on a 1.5T Signa Excite scanner. All CS reconstructions were implemented in Matlab (The MathWorks, Natick, MA) using the nonlinear conjugate gradient method as described in Appendix I. Two linear schemes were used for comparison, zero-filling with density compensation (ZF-w/dc) and low-resolution (LR). ZF-w/dc consists of a reconstruction by zero-filling the missing *k*-space data and *k*-space density compensation. The *k*-space density compensation is computed from the probability density function from which the random samples were drawn. LR consists of reconstruction from a Nyquist sampled low-resolution acquisition. The low-resolution acquisition contained centric-ordered data with the same number of data points as the undersampled sets. A software implementation of the reconstruction as well as some of the examples in this article are available at http://www.msrl.stanford.edu/~mlustig/software/

Simulation: CS Reconstruction Performance and Reconstruction Artifacts with Increased Undersampling

For the simulation we constructed a phantom by placing 18 features with 6 different sizes (3–75 pixel area) and 3 different intensities (0.33, 0.66, and 1). The features were distributed randomly in the phantom to simulate an angiogram. The phantom had 100×100 pixels out of which 575 are nonzero (5.75%). The finite-differences of the phantom consisted of 425 nonzeros (4.25%).

The first aim of the simulation was to examine the performance of the CS reconstruction and its associated artifacts with increased undersampling compared to the LR and ZF-w/dc methods. The second aim was to demonstrate the advantage of variable density random undersampling over uniform density random undersampling.

From the full *k*-space we constructed sets of randomly undersampled data with uniform density as well as variable density (density power of 12) with corresponding accelerations factors of 8, 12, and 20 (1,250, 834, and 500 *k*-space samples). Since the phantom is sparse both in image space and by finite differences, the data were CS reconstructed by using an ℓ^1 penalty on the image as well as a TV penalty (finite differences as the sparsifying transform) in Eq. [3]. The result was compared to the ZF-w/dc and LR linear reconstructions.

Undersampled 2D Cartesian Sampling in the Presence of Noise

CS reconstruction is known to be stable in the presence of noise (36,37), and can also be used to further perform nonlinear edge preserving denoising (29,38) of the image. To document the performance of CS in the presence of noise, we scanned a phantom using a 2D Cartesian spin-echo sequence with scan parameters yielding measured SNR = 6.17. The *k*-space was undersampled by a factor of 2.5 by randomly choosing phase-encodes lines with a quadratic variable density. A CS reconstruction using a TV penalty in Eq. [3] was obtained, with two different consistency RMS errors of $\epsilon = 10^{-5}$ and $\epsilon = 0.1$. The result was compared to the ZF-w/dc reconstruction, and the reconstruction based

on complete Nyquist sampling. Finally, the image quality as well as the resulting SNR of the reconstructions were compared.

Multislice 2DFT Fast Spin-Echo Brain Imaging

In the theory section, it was shown that brain images exhibit transform sparsity in the wavelet domain. Brain scans are a significant portion of MRI scans in the clinic, and most of these are multislice acquisitions. CS has the potential to reduce the acquisition time, or improve the resolution of current imagery.

In this experiment we acquired a T2-weighted multislice k-space data of a brain of a healthy volunteer using a FSE sequence ($256 \times 192 \times 32$, res = 0.82 mm, slice = 3 mm, echo-train = 15, TR/TE = 4,200/85 ms). For each slice we acquired different sets of 80 phase-encodes chosen randomly with quadratic variable density from 192 possible phase encodes, for an acceleration factor of 2.4. The image was CS reconstructed by using a wavelet transform (Daubechies 4) as sparsifying transform together with a TV penalty in Eq. [3]. To reduce computation time and memory load, we separated the 3D problem into many 2D CS reconstructions, i.e, iterating between solving for the $y - z$ plane slices, and solving for the $x - y$ plane slices. To demonstrate the reduction in scan time, as well as improved resolution, the multislice reconstruction was then compared to the ZF-w/dc and LR linear reconstructions and to the reconstruction based on complete Nyquist sampling.

The TPSF analysis shows that the multislice approach has considerable advantage over the 2DFT in recovering coarse scale image components. To demonstrate this, the multislice CS reconstruction was compared to a reconstruction from data in which each slice was undersampled in the same way. To further enhance the effect, we repeated the reconstructions for data that was randomly undersampled with uniform density where the coarse scale image components are severely undersampled.

Contrast-Enhanced 3D Angiography

Angiography is a very promising application for CS. First, the problem matches the assumptions of CS. Angiograms appear to be sparse already to the naked eye. The blood vessels are bright with a very low background signal. Angiograms are sparsified very well by both the wavelet transform and by finite-differences. This is illustrated in Fig. 3; blood vessel information is preserved in reconstructions using only 5% of the transform coefficients. Second, the benefits of CS are of real interest in this application. In angiography there is often a need to cover a very large FOV with relatively high resolution, and the scan time is often crucial.

To test the behavior of CS for various degrees of undersampling in a controlled way, we simulated k-space data by computing the Fourier transform of a magnitude post-contrast 3DFT angiogram of the peripheral legs. The scan was RF-spoiled gradient echo (SPGR) sequence with the following parameters: TR = 6 ms, TE = 1.5 ms, Flip = 30°. The acquisition matrix was set to $480 \times 480 \times 92$ with corresponding resolution of 1 mm $\times$ 0.8 mm $\times$ 1 mm. The imaging plane was coronal with a superior-inferior readout direction.

From the full k-space set, five undersampled data sets with corresponding acceleration factors of 5, 6.7, 8, 10, and 20 were constructed by randomly choosing phase encode lines with the quadratic variable k-space density. To reduce complexity, prior to reconstruction, a 1D Fourier transform was applied in the fully sampled readout direction. This effectively creates 480 separable purely random undersampled 2D reconstructions. Finally, the images were CS reconstructed by using a TV penalty in Eq. [3]. The result was compared to the ZF-w/dc and LR linear reconstructions.

We further tested the method, now with true k-space data on a first-pass abdominal contrast enhanced angiogram with the following scan parameters: TR/TE = 3.7/0.96 ms, FOV = 44 cm, matrix = $320 \times 192 \times 32$ (with 0.625 fractional echo), BW = 125 kHz.

The fully sampled data were undersampled 5-fold in retrospect with a quadratic k-space density effectively reducing the scan time from 22 s to 4.4 s. The images were CS reconstructed from the undersampled data using a TV penalty in Eq. [3] and the result was again compared to

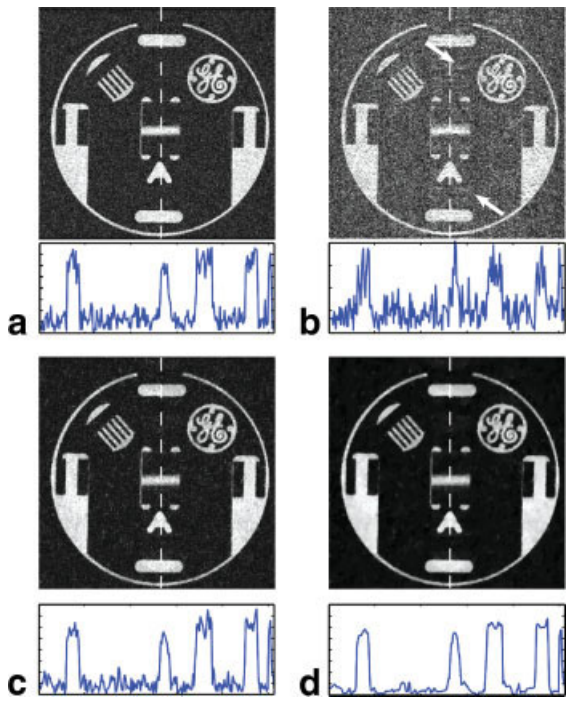

FIG. 7. 2DFT CS reconstruction from noisy data. CS reconstruction can perform denoising of the image as well as interference removal by relaxing the data consistency. (a) Reconstruction from complete noisy data. (b) ZF-w/dc, the image suffers from apparent noise due to incoherent aliasing as well as noise. (c) CS reconstruction with TV penalty from noisy undersampled data. Consistency RMS error set to 10^{-5}. (d) CS reconstruction with TV penalty from noisy undersampled data. Consistency RMS error set to 0.1. Note interference removal in both (c) and (d) and the denoising in (d). [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

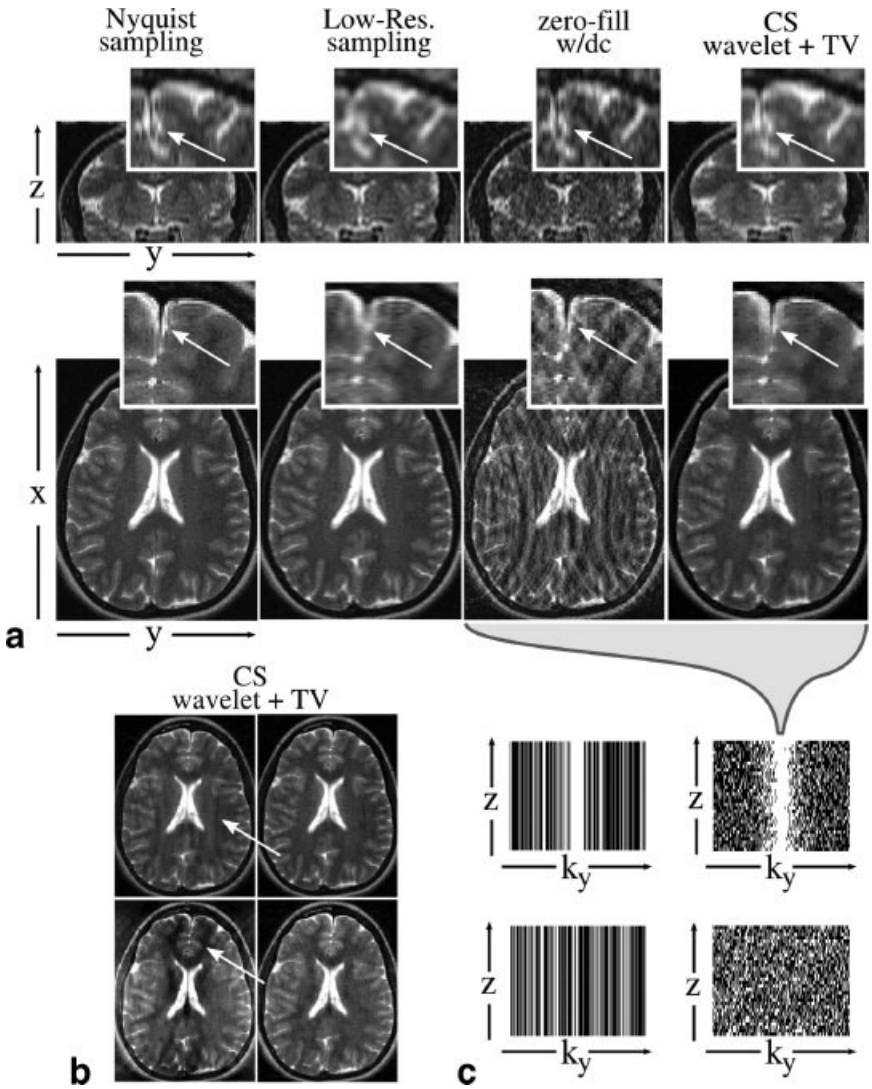

FIG. 8. Multislice 2DFT fast spin echo CS at 2.4 acceleration. (a) The CS-wavelet reconstruction exhibits significant resolution improvement over LR and significant suppression of the aliasing artifacts over ZF-w/dc compared to the full Nyquist sampling. (b) CS wavelet reconstructions from several undersampling schemes. The multi-slice approach outperforms the single-slice approach and variable density undersampling outperforms uniform undersampling. (c) The associated undersampling schemes; variable density (top) and uniform density (bottom), single-slice (left) and multi-slice (right).

the ZF-w/dc and LR linear reconstructions. To compensate for the fractional echo, a Homodyne partial-Fourier reconstruction (6) was performed in the readout direction.

RESULTS

Simulation: CS Reconstruction Performance and Reconstruction Artifacts with Increased Undersampling

Figure 6 presents the simulation results. The LR reconstruction, as expected, shows a decrease in resolution with acceleration characterized by loss of small structures and diffused boundaries. The ZF-w/dc reconstructions exhibit a decrease in apparent SNR because of the incoherent interference, which completely obscures small and dim features. The uniform density undersampling interference is significantly larger and more structured than the variable density. In both ZF-w/dc reconstructions the features that are brighter than the interference appear to have well-defined boundaries. In the CS reconstructions, at 8-fold acceleration (approximately 3 times more Fourier samples than sparse coefficients) we get exact recovery from both uniform density and variable density undersampling! At 12-fold acceleration (approximately 2 times more Fourier samples than sparse coefficients) we still get exact recovery from the variable density undersampling, but lose some of the low-contrast features in the uniform density undersampling. At 20-fold acceleration (similar number of Fourier samples as sparse coefficients) we get loss of image features in both reconstructions. The reconstruction errors are severe from the uniform density undersampling. However, in reconstruction from the variable density undersampling, only the weak intensity objects have reconstruction errors; the bright, high contrast features are well reconstructed.

2DFT CS Reconstruction in the Presence of Noise

Figure 7 presents the reconstruction results. Figure 7a shows the reconstruction of a fully sampled phantom scan. The measured SNR is 6.17. The ZF-w/dc reconstruction result in Fig. 7b exhibits significant apparent noise in the image with measured SNR of 3.79. The apparent noise is mostly incoherent aliasing artifacts due to the

undersampling as well as noise increase from the density compensation (which is essential to preserve the resolution). Some coherent aliasing artifacts are also visible (pointed to by arrows). In Fig. 7c the artifacts are suppressed by the CS reconstruction, recovering the noisy image with an SNR of 9.84. The SNR is slightly better because the CS reconstruction is inherently a denoising procedure. By increasing the RMS consistency parameter to $\epsilon = 0.1$ (less consistency) the CS reconstruction recovers and denoises the phantom image. Measured SNR increases dramatically to 26.9 without damaging the image quality. The denoising is nonlinear edge-preserving TV denoising and is shown in Fig. 7d.

Multislice Fast Spin-Echo Brain Imaging

Figure 8 shows the experiment results. In Fig. 8a coronal and axial slices of the multislice CS reconstruction are compared to the full Nyquist sampling, ZF-w/dc, and LR reconstructions. CS exhibits significant resolution improvement over LR and significant suppression of the aliasing artifacts over ZF-w/dc compared to the full Nyquist sampling.

Figure 8b shows CS reconstructions from several undersampling schemes. The corresponding undersampling schemes are given in Fig. 8c. Low-resolution aliasing artifacts are observed in the reconstructions in which the data was undersampled the same way for all slices. The artifacts are more pronounced for uniform undersampling. The reason is that some of the coarse-scale wavelet components in these reconstructions were not recovered correctly because of the large peak interference of coarse-scale components that was documented in the TPSF theoretical analysis (see Fig. 5a). These artifacts are significantly reduced when each slice is undersampled differently. This is because the theoretical TPSF peak interference in such sampling scheme is significantly smaller (see Fig. 5b), which enables better recovery of these components. The results in Fig. 8b show again that a variable density undersampling scheme performs significantly better than uniform undersampling.

Contrast Enhanced 3D Angiography

Figure 9 shows a region of interest in the maximum intensity projection (MIP) of the reconstruction results as well as a slice reconstruction from 10-fold acceleration. The LR reconstruction (left column), as expected, shows a decrease in resolution with acceleration characterized by loss of small structures and diffused blood vessel boundaries. The ZF-w/dc reconstruction (middle column), exhibits a decrease in apparent SNR because of the incoherent interference, which obscures small and dim vessels. Interestingly, the boundaries of the very bright vessels remain sharp and are diagnostically more useful than the LR. The CS reconstruction (right column), on the other hand, exhibits good reconstruction of the blood vessels even at very high accelerations. The resolution as well as the contrast are preserved with almost no loss of information at up to 10-fold acceleration. Even at acceleration of 20-fold the bright blood vessel information is well preserved. These results conform with the thresholding experiment in Fig. 3 as well as the simulation results in Fig. 6.

Figure 10 shows the reconstruction result from the first-pass contrast experiment. The imaged patient has an aorto-bifemoral bypass graft. This is meant to carry blood from the aorta to the lower extremities and is seen on the left side of the aorta (right in the image). There is a high-grade stenosis in the native right common illiac artery, which is indicated by the arrows. Again, at 5-fold acceleration the LR acquisition exhibits diffused boundaries and the ZF-w/dc exhibits considerable decrease in apparent SNR. The CS reconstruction exhibits a good reconstruction of the blood vessels, in particular, we see that in Fig. 10d flow across the stenosis is visible, but it is not visible in Figs. 10b,c.

DISCUSSION

Computational Complexity

Development of fast algorithms for solving Eq. [3] accurately or approximately is an increasingly popular research topic. Many of these methods have been mentioned in the Theory section. Overall, the reconstruction is iterative and more computationally intensive than linear reconstruction methods. However, some of the methods proposed show great potential to significantly reduce the overall complexity.

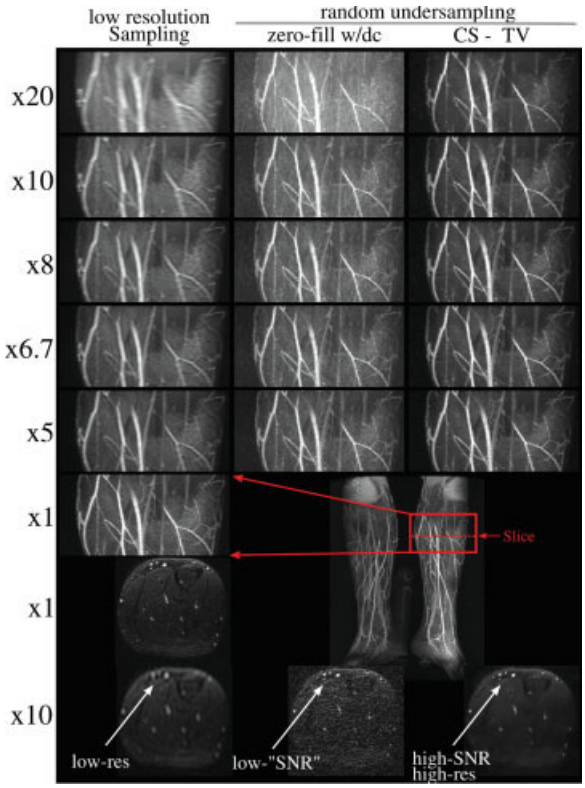

FIG. 9. Contrast-enhanced 3D angiography reconstruction results as a function of acceleration. Left column: acceleration by LR. Note the diffused boundaries with acceleration. Middle column: ZF-w/dc reconstruction. Note the increase of apparent noise with acceleration. Right column: CS reconstruction with TV penalty from randomly undersampled *k*-space. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

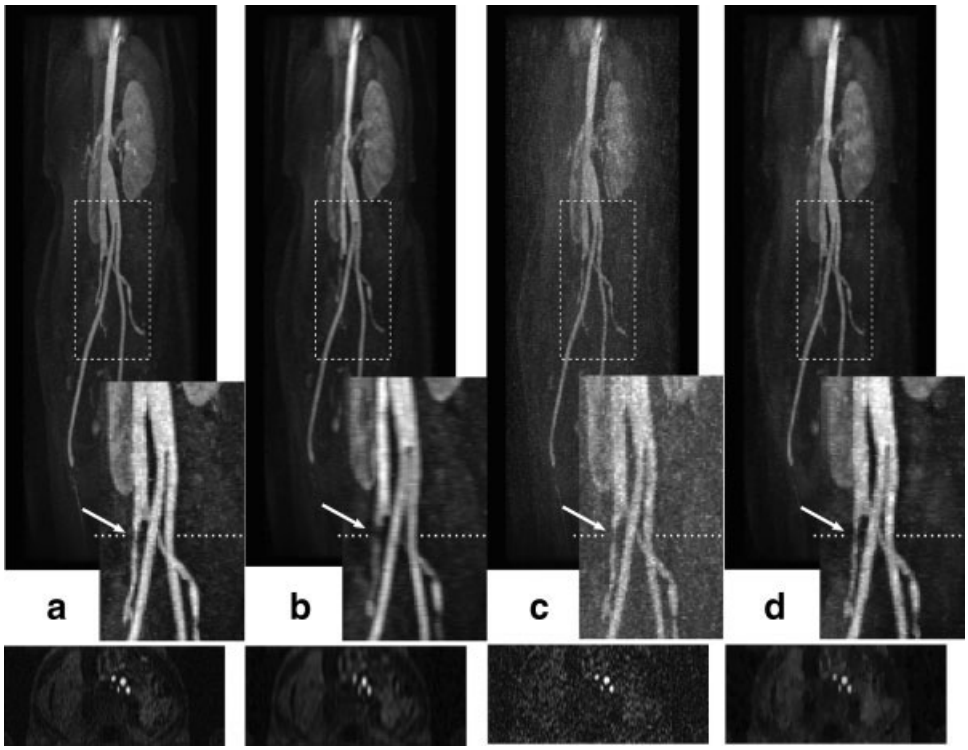

FIG. 10. Reconstruction from 5-fold accelerated acquisition of first-pass contrast enhanced abdominal angiography. (a) Reconstruction from a complete data set. (b) LR (c) ZF-w/dc (d) CS reconstruction from random undersampling. The patient has a aorto-bifemoral bypass graft. This is meant to carry blood from the aorta to the lower extremities. There is a high-grade stenosis in the native right common illiac artery, which is indicated by the arrows. In figure parts (a) and (d) flow across the stenosis is visible, but it is not on (b) and (c).

The examples in this article were reconstructed using a nonlinear conjugate gradient method with backtracking line-search. In a Matlab (The MathWorks, Natick, MA) implementation, it takes about 150 CG iterations (approximately 30 s) to reconstruct a 480×92 angiogram using a TV-penalty at 5-fold acceleration. We expect a significant reduction in the reconstruction time by code optimization.

Reconstruction Artifacts

The ℓ^1 reconstruction tends to slightly shrink the magnitude of the reconstructed sparse coefficients. The resulting reconstructed coefficients are often slightly smaller than in the original signal. This coefficient shrinkage decreases when the reconstruction consistency parameter ϵ in Eq. [3] is small.

In some wavelet-based CS reconstructions, small high-frequency oscillatory artifacts may appear in the reconstruction. This is due to false detection of fine-scale wavelet components. To mitigate these artifacts it is recommended to add a small TV penalty on top of the wavelet penalty. This can be considered as requiring the image to be sparse in both wavelet and finite-differences transforms.

In CS, the contrast in the image plays a major part in the ability to vastly undersample and reconstruct images. High contrast often results in large distinct sparse coefficients. These can be recovered even at very high accelerations. For example, a single bright pixel will most likely appear in the reconstruction even with vast undersampling (See Figs. 6 and 9 for an example). However, features with lower contrast at the same accelerations will be so deeply submerged by the interference that they would not be recoverable. As such, with increased acceleration the most distinct artifacts in CS are not the usual loss of resolution or increase in aliasing interference, but loss of low-contrast features in the image. Therefore, CS is particularly attractive in applications that exhibit high resolution high contrast image features, and rapid imaging is required.

Relation to Other Acceleration Methods

Vastly undersampled 3D radial trajectories–VIPR (39) have demonstrated high acceleration for angiography. The VIPR trajectory is a 3D incoherent sampling scheme in which the interference spreads in all three dimensions. As such, reconstruction from VIPR acquisitions can be further improved by using the CS approach.

Wajer's PhD thesis (11) suggested undersampling k-space and employing a Bayesian reconstruction to randomized trajectories. This approach, although different, is related to finite difference sparsity.

Nonuniform sampling with maximum entropy reconstruction has been used successfully to accelerate multi-dimensional NMR acquisitions (40). Maximum entropy reconstruction is also related to sparsity of finite differences.

CS reconstruction exploits sparsity and compressibility of MR images. It can be combined with other acceleration methods that exploit different redundancies. For example, constraining the image to be real in Eq. [4] effectively combines phase constrained partial k-space with the CS reconstruction. In a similar way, CS can be combined with SENSE reconstruction by including the coil sensitivity information in Eq. [3]. In general, any other prior on the image that can be expressed as a convex constraint can be incorporated in the reconstruction.

CONCLUSIONS

We have presented the theory of CS and the details of its implementation for rapid MR imaging. We demonstrated experimental verification of several implementations for 2D and 3D Cartesian imaging. We showed that the sparsity of MR images can be exploited to significantly reduce scan time, or alternatively, improve the resolution of MR imagery. We demonstrated high acceleration in in-vivo experiments, in particular a 5-fold acceleration of first pass contrast enhanced MRA. CS can play a major part in many applications that are limited by the scan time, when the images exhibit transform sparsity.

ACKNOWLEDGMENTS

The authors would like to thank Walter Bloch for his help in the project and Marcus Alley for providing some of the experimental data. The authors would like to thank Peder Larson, Brian Hargreaves, William Overall, Nikola Stikov, and Juan Santos for their comments and help in the preparation of this manuscript.

APPENDIX A: NONLINEAR CONJUGATE-GRADIENT SOLUTION OF THE CS OPTIMIZATION PROCEDURE

Equation [3] poses a constrained convex optimization problem. Consider the unconstrained problem in so-called Lagrangian form:

$$\underset{\boldsymbol{m}}{\operatorname{argmin}} \quad \|\mathcal{F}_u m - y\|_2{}^2 + \lambda \|\Psi m\|_1, \qquad [A1]$$

where λ is a regularization parameter that determines the trade-off between the data consistency and the sparsity. As is well-known, the parameter λ can be selected appropriately such that the solution of Eq. [A1] is exactly as Eq. [3]. The value of λ can be determined by solving Eq. [A1] for different values, and then choosing λ so that $||F_u m - y||_2 \approx \epsilon$.

We propose solving Eq. [A1] using a nonlinear conjugate gradient descent algorithm with backtracking line search where $f(m)$ is the cost-function as defined in Eq. [A1].

ITERATIVE ALGORITHM FOR ℓ^1-PENALIZED RECONSTRUCTION

INPUTS:
- y - k-space measurements
- $\mathcal{F}_u$ - undersampled Fourier operator associated with the measurements
- Ψ - sparsifying transform operator
- λ - a data consistency tuning constant

OPTIONAL PARAMETERS:
- TolGrad - stopping criteria by gradient magnitude (default 10^{-4})
- MaxIter - stopping criteria by number of iterations (default 100)
- α, β - line search parameters (defaults $\alpha = 0.05$, $\beta = 0.6$)

OUTPUTS:
- m - the numerical approximation to Eq. [A1]

% Initialization
$k = 0$; $m = 0$; $g_0 = \nabla f(m_0)$; $\Delta m_0 = -g_0$
% Iterations
while ($||g_k||_2 <$ TolGrad and $k >$ maxIter) {
  % Backtracking line-search
  $t = 1$; *while* ($f(m_k + t\Delta m_k) > f(m_k) + \alpha t \cdot Real(g_k^* \Delta m_k)$)
    $\{t = \beta t\}$
  $m_{k+1} = m_k + t\Delta m_k$
  $g_{k+1} = \nabla f(m_{k+1})$
  $\gamma = \frac{||g_{k+1}||_2^2}{||g_k||_2^2}$
  $\Delta m_{k+1} = -g_{k+1} + \gamma \Delta m_k$
  $k = k + 1$ }

The conjugate gradient requires the computation of $\nabla f(m)$ which is,

$$\nabla f(m) = 2F_u^*(F_u m - y) + \lambda \nabla ||\Psi m||_1 \qquad [A2]$$

The ℓ^1 norm is the sum of absolute values. The absolute value function, however, is not a smooth function and as a result Eq. [A2] is not well defined for all values of m. Instead, we approximate the absolute value with a smooth function by using the relation $|x| \approx \sqrt{x^* x + \mu}$, where μ is a positive smoothing parameter. With this approximation, $\frac{d|x|}{dx} \approx \frac{x}{\sqrt{x^* x + \mu}}$.

Now, let W be a diagonal matrix with the diagonal elements $w_i = \sqrt{(\Psi m)_i^* (\Psi m)_i + \mu}$. Equation [A2] can be approximated by,

$$\nabla f(m) \approx 2F_u^*(F_u m - y) + \lambda \Psi^* W^{-1} \Psi m \qquad [A3]$$

In practice, Eq. [A3] is used with a smoothing factor $\mu \in [10^{-15}, 10^{-6}]$. The number of CG iterations varies with different objects, problem size, accuracy and undersampling. Examples in this paper required between 80 and 200 CG iterations.

APPENDIX B: DERIVATION OF THE INTERFERENCE STANDARD DEVIATION FORMULA

Equation [1] is easily derived. The total energy in the PSF is $\frac{N}{D}$ and the energy of the main lobe is $\left(\frac{N}{D}\right)^2$. The off-center energy is therefore $\frac{N}{D} - \left(\frac{N}{D}\right)^2$. Normalizing by the number of off-center pixels and also by the main lobe's energy and setting $p = \frac{D}{N}$ we get Eq. [1].

REFERENCES

1. Marseille GJ, de Beer R, Fuderer M, Mehlkopf AF, van Ormondt D. Nonuniform phase-encode distributions for MRI scan time reduction. J Magn Reson 1996;111:70–75.
2. Scheffler K, Hennig J. Reduced circular field-of-view imaging. Magn Reson Med 1998;40:474–480.
3. Tsai CM, Nishimura D. Reduced aliasing artifacts using variable-density k-space sampling trajectories. Magn Reson Med 2000;43:452–458.
4. 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.
5. Greiser A, von Kienlin M. Efficient k-space sampling by density-weighted phase-encoding. Magn Reson Med 2003;50:1266–1275.
6. 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.

7. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: Sensitivity encoding for fast MRI. Magn Reson Med 1999;42:952–962.
8. Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays. Magn Reson Med 1997;38:591–603.
9. Korosec FR, Frayne R, Grist TM, Mistretta CA. Time-resolved contrast-enhanced 3D MR angiography. Magn Reson Med 1996;36:345–351.
10. Madore B, Glover G, Pelc N. Unaliasing by fourier-encoding the overlaps using the temporal dimension (UNFOLD), applied to cardiac imaging and fMRI. Magn Reson Med 1999;42:813–828.
11. Wajer F. Non-cartesian MRI scan time reduction through sparse sampling. PhD thesis, Delft University of Technology, 2001.
12. 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.
13. Mistretta CA, Wieben O, Velikina J, Block W, Perry J, Wu Y, Johnson K, Wu Y. Highly constrained backprojection for time-resolved MRI. Magn Reson Med 2006;55:30–40.
14. Candès E, Romberg J, Tao T. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. IEEE Trans Inf Theory 2006;52:489–509.
15. Donoho D. Compressed sensing. IEEE Trans Inf Theory 2006;52:1289–1306.
16. Taubman DS, Marcellin MW. JPEG 2000: Image compression fundamentals, standards and practice. Kluwer International Series in Engineering and Computer Science; 2002. Kluwer Academic Publishers.
17. Lustig M, Lee JH, Donoho DL, Pauly JM. Faster imaging with randomly perturbed, under-sampled spirals and ℓ^1 reconstruction. In Proceedings of the 13th Annual Meeting of ISMRM, Miami Beach, 2005. p 685.
18. Santos JM, Cunningham CH, Lustig M, Hargreaves BA, Hu BS, Nishimura DG, Pauly JM. Single breath-hold whole-heart MRA using variable-density spirals at 3T. Magn Reson Med 2006;55:371–379.
19. Chang TC, He L, Fang T. MR image reconstruction from sparse radial samples using bregman iteration. In Proceedings of the 13th Annual Meeting of ISMRM, Seattle, 2006. p 696.
20. Ye JC, Tak S, Han Y, Park HW. Projection reconstruction MR imaging using FOCUSS. Magn Reson Med 2007;57:764–775.
21. 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.
22. Fain SB, Block W, Charles A, Mistretta AB. Correction for artifacts in 3D angularly undersampled MR projection reconstruction. In Proceedings of the 9th Annual Meeting of ISMRM, Glasgow, 2001. p 759.
23. Starck J, Elad M, Donoho D. Image decomposition via the combination of sparse representations and a variational approach. IEEE Trans Image Process 2005;14:1570–1582.
24. Lustig M, Santos JM, Donoho DL, Pauly JM. k-t SPARSE: High frame rate dynamic MRI exploiting spatio-temporal sparsity. In Proceedings of the 13th Annual Meeting of ISMRM, Seattle, 2006. p 2420.
25. Jung H, Ye JC, Kim EY. Improved k-t BLAST and k-t SENSE using FOCUSS. Phys Med Biol 2007;52:3201–3226.
26. Candés E, Romberg JK. Signal recovery from random projections. In Proceedings of SPIE Computational Imaging III, San Jose, 2005. p 5674.
27. Tsaig Y, Donoho DL. Extensions of compressed sensing. Signal Process 2006;86:533–548.
28. Chen S, Donoho D, Saunders M. Atomic decomposition by basis pursuit. SIAM J Sci Comput 1999;20:33–61.
29. Rudin L, Osher S, Fatemi E. Non-linear total variation noise removal algorithm. Phys D 1992;60:259–268.
30. Kim SJ, Koh K, Lustig M, Boyd S. An efficient method for compressed sensing. In Proceedings of IEEE International Conference on Image Processing (ICIP), San Antonio, 2007, in press.
31. Daubechies I, Defrise M, Mol CD. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. Commun Pure Appl Math 2004;57:1413–1457.
32. Starck JL, Elad M, Donoho D. Image decomposition via the combination of sparse representation and a variational approach. IEEE Trans Image Process 2005;14:1570–1582.
33. Elad M, Matalon B, Zibulevsky M. Coordinate and subspace optimization methods for linear least squares with non-quadratic regularization. J Appl Comput Harmonic Anal 2006, doi: 10.1016/j.acha.2007.02.002.
34. Donoho D, Elad M, Temlyakov V. Stable recovery of sparse overcomplete representations in the presence of noise. IEEE Trans Inf Theory 2006;52:6–18.
35. Bronstein MM, Bronstein AM, Zibulevsky M, Azhari H. Reconstruction in diffraction ultrasound tomography using nonuniform FFT. IEEE Trans Med Imaging 2002;21:1395–1401.
36. Candés E, Romberg J, Tao T. Stable signal recovery from incomplete and inaccurate measurements. Commun Pure Appl Math 2006;59:1207–1223.
37. Haupt J, Nowak R. Signal reconstruction from noisy random projections. IEEE Trans Inf Theory 2006;52:4036–4048.
38. Donoho D, Johnstone I. Ideal spatial adaptation via wavelet shrinkage. Biometrika 1994;81:425–455.
39. Barger AV, Bloch 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.
40. Rovnyak D, Frueh D, Sastry M, Sun Z, Stern A, Hoch J, Wagner G. Accelerated acquisition of high resolution triple-resonance spectra using non-uniform sampling and maximum entropy reconstruction. J Magn Reson 2004;170:15–21.