



Compressive
Sampling

© DIGITAL VISION

Compressed Sensing MRI

[A look at how CS can improve on current imaging techniques]

[Michael Lustig,
David L. Donoho,
Juan M. Santos, and
John M. Pauly]

Compressed sensing (CS) aims to reconstruct signals and images from significantly fewer measurements than were traditionally thought necessary. Magnetic resonance imaging (MRI) is an essential medical imaging tool with an inherently slow data acquisition process. Applying CS to MRI offers potentially significant scan time reductions, with benefits for patients and health care economics.

MRI obeys two key requirements for successful application of CS: 1) medical imagery is naturally *compressible* by sparse coding in an appropriate transform domain (e.g., by wavelet transform), and 2) MRI scanners naturally acquire *encoded* samples, rather than direct pixel samples (e.g., in *spatial-frequency encoding*).

In this article we review the requirements for successful CS, describe their natural fit to MRI, and then give examples of four interesting applications of CS in MRI. We emphasize

Digital Object Identifier 10.1109/MSP.2007.914728

an intuitive understanding of CS by describing the CS reconstruction as a process of interference cancellation. We also emphasize an understanding of the driving factors in applications, including limitations imposed by MRI hardware, by the characteristics of different types of images, and by clinical concerns.

PRINCIPLES OF MAGNETIC RESONANCE IMAGING

We first briefly sketch properties of MRI related to CS. More complete descriptions of MRI can be found in the excellent survey paper by Wright [1] from this magazine and in MRI textbooks.

NUCLEAR MAGNETIC RESONANCE PHYSICS

The MRI signal is generated by protons in the body, mostly those in water molecules. A strong static field B_0 polarizes the protons, yielding a net magnetic moment oriented parallel to the static field. Applying a radio frequency (RF) excitation field B_1 produces a magnetization component m transverse to the static field. This magnetization precesses at a frequency proportional to the static field strength. The transverse component of the precessing magnetization emits a radio frequency signal detectable by a receiver coil. The transverse magnetization $m(\vec{r})$ at position \vec{r} and its corresponding emitted RF signal can be made proportional to many different physical properties of tissue. One property is the proton density, but other properties [1] can be emphasized as well. MR image reconstruction attempts to visualize $m(\vec{r})$, depicting the spatial distribution of the transverse magnetization.

SPATIAL ENCODING

MR systems encode spatial information in the MR signal by superimposing additional magnetic fields on top of the strong static field. These so-called gradient fields vary linearly in space and are denoted as G_x , G_y , and G_z corresponding to the three Cartesian axes. When G_x is applied, the magnetic field will vary with position as $B(x) = B_0 + G_x x$, causing the precession frequency to vary linearly in space. As a result, magnetization at positive x positions will precess at a higher frequency than magnetization at negative x positions.

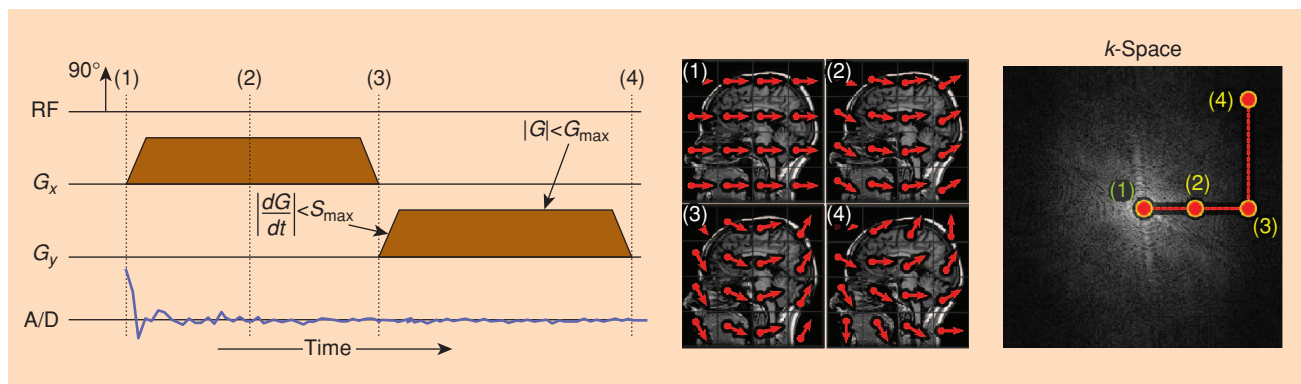
Spatial encoding using gradients can be understood by analogy with the piano. The pitch of a piano note varies linearly with the position of the key being struck; the sound one hears is the net sum of all notes emitted. A skilled musician listening to the emitted polyharmonic sound can hear which notes are playing and say which keys were pressed, and how forcefully. The MR signal generated in the presence of a gradient field is likewise a polyphonic mixture. The spatial positions within the patient's body are like piano keys. The emitted RF signal from each position is like a "note," with a frequency linearly varying with position. The polyharmonic MR signal superimposes the different "notes"; they encode the spatial position and the magnetization strength at those positions. A signal processing engineer will recognize the Fourier relation between the received MR signal and the magnetization distribution and that the magnetization distribution can be decoded by a spectral decomposition.

Multidimensional spatial encoding can be further understood by introducing the notion of k -space. Gradient-induced variation in precession frequency causes a location-dependent linear phase dispersion to develop. Therefore the receiver coil detects a signal encoded by the linear phase. It can be shown [1] that the signal equation in MRI has the form of a Fourier integral,

$$s(t) = \int_R m(\vec{r}) e^{-i2\pi \vec{k}(t) \cdot \vec{r}} d\vec{r},$$

where $\vec{k}(t) \propto \int_0^t \vec{G}(s) ds$. In words, the received signal at time t is the Fourier transform of the object $m(\vec{r})$ sampled at the spatial frequency $\vec{k}(t)$. Such Fourier encoding is fundamentally encoded and very different than traditional optical imaging where pixel samples are measured directly.

The design of an MRI acquisition method centers on developing the gradient waveforms $\vec{G}(t) = [G_x(t), G_y(t), G_z(t)]^T$ that drive the MR system. These waveforms, along with the associated RF pulses used to produce the magnetization, are called a *pulse sequence*. The integral of the $\vec{G}(t)$ waveforms traces out a trajectory $\vec{k}(t)$ in spatial frequency space, or k -space. For illustration, consider the simple example in Figure 1



[FIG1] The temporal MRI signal directly samples the spatial frequency domain of the image. Gradient fields cause a linear frequency distribution across the image, which produces a linear phase accrual with time. The received signals are spatial frequency samples of the image. The corresponding spatial frequencies are proportional to the gradient waveform area. The gradient is limited in amplitude, G_{\max} , and slew rate, S_{\max} , which are both system specific.

where, immediately after the RF excitation, a G_x gradient field is applied followed by a G_y gradient. The phases of the magnetization are shown at different time points, along with the k -space trajectory and the MR signal. This encoded sampling and the freedom in choosing the sampling trajectory play a major role in making CS ideas naturally applicable to MRI.

IMAGE ACQUISITION

Constructing a single MR image commonly involves collecting a series of frames of data, called *acquisitions*. In each acquisition, an RF excitation produces new transverse magnetization, which is then sampled along a particular trajectory in k -space. Due to various physical and physiological constraints [1], most MRI imaging methods use a sequence of acquisitions and each one samples part of k -space. The data from this sequence of acquisitions are then used to reconstruct an image. Traditionally the k -space sampling pattern is designed to meet the Nyquist criterion, which depends on the resolution and field of view (FOV) as shown in Figure 2. Violation of the Nyquist criterion causes image artifacts in linear reconstructions. The appearance of such artifacts depends on the details in the sampling pattern, as discussed below.

In MRI, it is possible to selectively excite a thin slice through the 3-D volume. This reduces the data collection to two dimensions in k -space for each slice. The volumetric object is imaged by exciting more slices, known as a multislice acquisition. When a volume or a thick slab is excited, a 3-D k -space volume must be sampled. Each of these approaches is very common and has advantages in specific applications.

We have considerable freedom in designing the k -space trajectory for each acquisition. Some trajectories are illustrated in Figure 2. By far the most popular trajectory uses

straight lines from a Cartesian grid. Most pulse sequences used in clinical imaging today are Cartesian. Reconstruction from such acquisitions is wonderfully simple: apply the inverse fast Fourier transform (FFT). More importantly, reconstructions from Cartesian sampling are robust to many sources of system imperfections.

While Cartesian trajectories are by far the most popular, some other trajectories are in use, including sampling along radial lines and sampling along spiral trajectories. Radial acquisitions are less susceptible to motion artifacts than Cartesian trajectories and can be significantly undersampled [2], especially for high contrast objects [3]. Spirals make efficient use of the gradient system hardware and are used in real-time and rapid imaging applications [4]. Efficient reconstruction from such non-Cartesian trajectories requires using filtered back-projection or interpolation schemes (e.g., gridding [5]).

RAPID IMAGING


MR acquisition is inherently a process of traversing curves in multidimensional k -space. The speed of k -space traversal is limited by physical constraints. In current systems, gradients are limited by maximum amplitude and maximum slew-rate (see Figure 1). In addition, high gradient amplitudes and rapid switching can produce peripheral nerve stimulation [1]. Since this must be avoided, physiology provides a fundamental limit to gradient system performance.

This fundamental limit has caused many researchers to search for methods to reduce the amount of acquired data without degrading image quality. Many such efforts are inspired by the idea that MRI data are redundant. Such redundancy can be created by design. For example, using multiple receiver coils [6], [7] provides more useful data per MR acquisition, requiring fewer acquisitions per scan. Redundancy can be a known or modeled signal property such as spatial-temporal correlation [8]–[11] or a redundancy learned and extracted from the data itself [12]–[14].

All efforts at reduced data acquisition might well be labeled “compressive sampling,” however, the underlying phenomena being exploited are often quite different. In this article, we focus on approaches rooted in the theory described in [15]–[17]; such approaches are called here CS approaches. Much ongoing work is based on such approaches [18]–[23].

THE SPARSITY/COMPRESSIBILITY OF MR IMAGES

Natural images can often be compressed with little or no perceptible loss of information [24]. The world-wide-Web demonstrates this billions of times weekly. Transform-based compression is a widely used strategy adopted in the JPEG, JPEG-2000, and MPEG standards. This strategy first applies a sparsifying transform, mapping image content into a vector of sparse coefficients, and then encodes the sparse vector by approximating the most significant coefficients and ignoring the smaller ones. The discrete wavelet transform (DWT) is a common sparsifying transform and is at the heart of JPEG-2000 [24].



[FIG2] The Nyquist criterion sets the required k -space coverage, which can be achieved using various sampling trajectories. Image resolution is determined by the extent of the k -space coverage. The supported field of view is determined by the sampling density. Violation of the Nyquist criterion causes artifacts in linear reconstructions, which depend on the sampling pattern.

Most MR images are sparse in an appropriate transform domain. To begin with, consider Figure 3. Angiograms, which are images of blood vessels, contain primarily contrast-enhanced blood vessels in a sea of void and already look sparse in the pixel representation. They can be made even sparser by spatial finite-differencing. More complex imagery, such as brain images, can be sparsified in more sophisticated domains, such as the wavelet domain. Sparse representation is not limited to still imagery. Often videos can safely be compressed much more heavily. This is demonstrated by the success of MPEG. Dynamic MR images are highly compressible as well. For example, the quasi-periodicity of heart images has a sparse temporal Fourier transform.

THE NATURAL FIT BETWEEN CS AND MRI

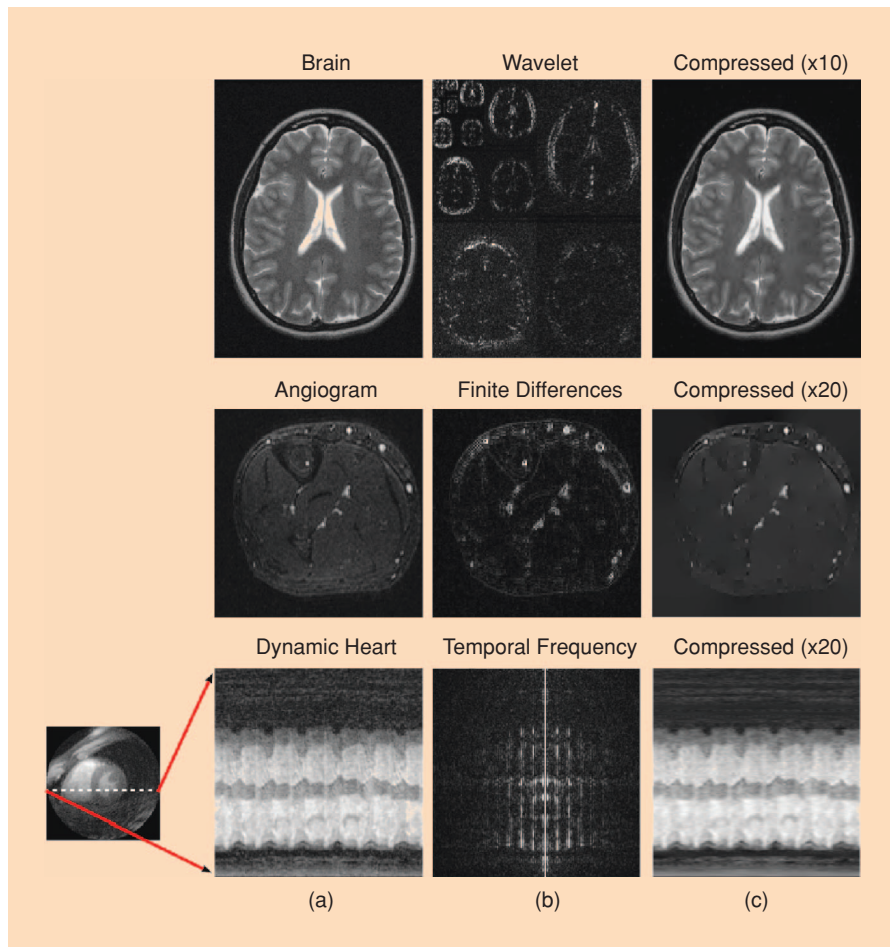
The transform sparsity of MR images and the coded nature of MR acquisition are two key properties enabling CS in MRI. Figure 4 illustrates these elements, making MRI a natural CS system. We now give a more formal discussion of the requirements.

COMPRESSED SENSING THEORY

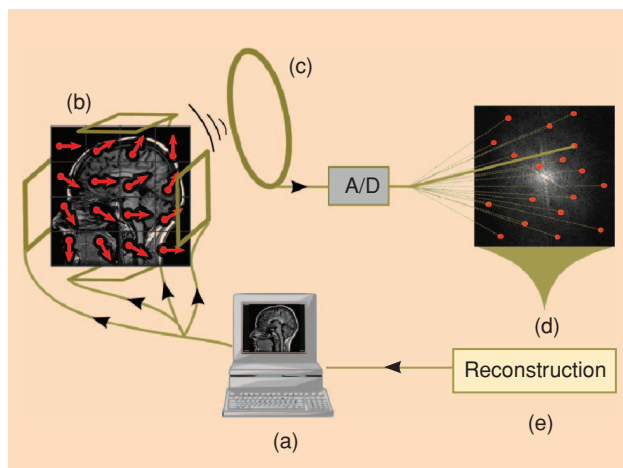
CS emerged in the literature of information theory and approximation theory as an abstract mathematical idea [15]–[17]. One measures a relatively small number of “random” linear combinations of the signal values—much smaller than the number of signal samples nominally defining it. However, because the underlying signal is compressible, the nominal number of signal samples is a gross overestimate of the “effective” number of “degrees of freedom” of the signal. As a result, the signal can be reconstructed with good accuracy from relatively few 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 original paper by Candès, Romberg, and Tao [15] was motivated in large part by MRI since it looked at random undersampling of Fourier coefficients.

Theoretical and technical aspects of CS are discussed elsewhere in this special issue. However, the key points can be reduced to nontechnical language. A successful application of CS has three requirements:



[FIG3] Transform sparsity of MR images. (a) Fully sampled images are mapped by a sparsifying transform to a (b) transform domain; the several largest coefficients are preserved while all others are set to zero; the transform is inverted forming a (c) reconstructed image.



[FIG4] MRI as a compressed sensing system. The user controls the gradient and RF waveforms that, in turn, control the phase of the pixels/voxels in the image. An RF coil receives the signal in an encoded form—samples in k -space. Careful crafting of the gradient waveforms allows for incoherent measurements of k -space. With an appropriate nonlinear reconstruction enforcing sparsity, an image can be reconstructed.

- **Transform sparsity:** The desired image should have a sparse representation in a known transform domain (i.e., it must be compressible by transform coding).
- **Incoherence of undersampling artifacts:** The artifacts in linear reconstruction caused by k -space undersampling should be incoherent (noise like) in the sparsifying transform domain.
- **Nonlinear reconstruction:** The image should be reconstructed by a nonlinear method that enforces both sparsity of the image representation and consistency of the reconstruction with the acquired samples.

The first condition is clearly met for MR images, as explained above. The fact that incoherence is important, that MR acquisition can be designed to achieve incoherent undersampling, and the fact that there are efficient and practical algorithms for reconstruction will not, at this point in the article, be at all obvious. So we turn to a very simple example.

INTUITIVE EXAMPLE: INTERFERENCE CANCELLATION

To develop intuition for the importance of incoherence and the feasibility of CS, consider the one-dimensional (1-D) case illustrated in Figure 5. A sparse signal, Figure 5(a), is sub-Nyquist (eight-fold) sampled in k -space [Figure 5(b)]. Simply zero-filling the missing values and inverting the Fourier transform results in artifacts that depend on the sampling pattern. With equispaced undersampling [Figure 5(d)], this reconstruction generates a superposition of shifted signal copies. In this case, recovery of the original signal is hopeless, as each replica is an equally likely candidate.

With random undersampling, the situation is very different. The zero-filled Fourier reconstruction exhibits incoherent artifacts that actually behave much like additive random noise [Figure 5(c)]. Despite appearances, the artifacts are not noise; rather, undersampling causes leakage of energy away from each individual nonzero value of the original signal. This energy

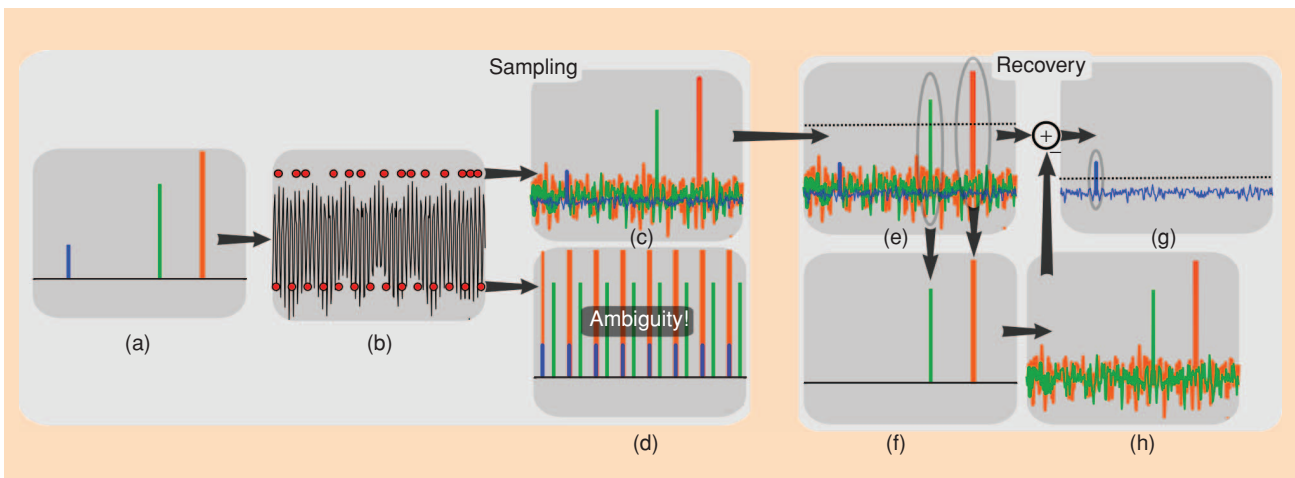
appears in other reconstructed signal coefficients, including those which had been zero in the original signal.

It is possible, starting from knowledge of the k -space sampling scheme and the underlying original signal, to calculate this leakage analytically. This observation immediately suggests a nonlinear iterative technique which enables accurate recovery, even though the signal in Figure 5(a) was eight-fold undersampled.

A simple heuristic recovery procedure is illustrated in Figure 5(e)–(h). It applies iterative thresholding, picking the largest components of the signal, calculating the interference that would be caused by the presence of those components and subtracting it. After subtracting the calculated interference, smaller components, previously submerged in interference, rise above it and can be recovered [25].

INCOHERENT SAMPLING IN MRI

Designing a CS scheme for MRI can now be viewed as selecting a subset of the frequency domain that can be efficiently sampled and is incoherent with respect to the sparsifying transform. Before we formally introduce the notion of incoherence, we note that narrow optimization of incoherence must not be pushed too far. Some of the most powerful and elegant results about CS assume one samples a completely random subset of k -space, which indeed gives very low coherence [15]. The motivation for random sampling can be easily and intuitively understood using our 1-D example given earlier. Although random sampling is an inspiring and instructive idea, sampling a truly random subset of k -space is generally impractical. Any practical sampling trajectory must satisfy hardware and physiological constraints. Therefore sampling trajectories must follow relatively smooth lines and curves. Sampling schemes must also be robust to nonideal, real-life situations. Non-Cartesian sampling schemes can be highly sensitive to system imperfections.



[FIG5] Heuristic procedure for reconstruction from undersampled data. A sparse signal (a) is 8-fold undersampled in its 1-D k -space domain (b). Equispaced undersampling results in signal aliasing (d) preventing recovery. Pseudo-random undersampling results in incoherent interference (c). Some strong signal components stick above the interference level, are detected and recovered by thresholding (e) and (f). The interference of these components is computed (g) and subtracted (h), thus lowering the total interference level and enabling recovery of weaker components.

Furthermore, a uniform random distribution of samples in spatial frequency does not take into account the energy distribution of MR images in k -space, which is far from uniform. Most energy in MR imagery is concentrated close to the center of k -space and rapidly decays towards the periphery of k -space.

Therefore, realistic designs for CS in MRI should have variable-density sampling with denser sampling near the center of k -space, matching the energy distribution in k -space. Such designs should also create k -space trajectories that are somewhat irregular and partially mimic the incoherence properties of pure random sampling, yet allow rapid collection of data. To compare designs, we need a quantitative notion of incoherence that will allow us to compare data acquisition schemes and their performance.

THE TRANSFORM SPARSITY OF MR IMAGES AND THE CODED NATURE OF MR ACQUISITION ARE TWO KEY PROPERTIES ENABLING CS IN MRI.

MEASURING INCOHERENCE

We first measure incoherence for cases where the image is already sparse in the pixel domain, so no further sparsification is needed. Suppose we sample a subset S of k -space. Let \mathcal{F}_S denote the Fourier transform evaluated just at frequencies in subset S . Let \mathcal{F}_S^* denote the adjoint operation, which can be represented as zero-filling followed by inverse Fourier transform. Define the *point spread function* (PSF) as, simply $\text{PSF}(i, j) = (\mathcal{F}_S^* \mathcal{F}_S)(i, j)$. Under complete Cartesian sampling, the PSF is the identity and off-diagonal terms vanish. Undersampling k -space induces nonzero off-diagonal terms in $\text{PSF}(i, j)$. A nonzero at (i, j) means that linear reconstruction of pixel i suffers interference by a unit impulse in pixel $j \neq i$. In short, the PSF measures the tendency of zero-filled linear reconstruction to leak energy from the true underlying source pixel to other pixels. This energy shows up as blurring or aliasing artifacts in the image. The goal of irregular sampling is to spread such leakage quasi-uniformly across the image, so that the maximal leakage is small. Accordingly, we define coherence as the maximum off-diagonal entry in a properly normalized PSF. This is analogous to notions of sidelobe-to-peak ratio the reader will have encountered in many branches of signal processing. Figure 6 shows some PSFs of irregular trajectories.

Most MR images are sparse in a transform domain other than the pixel domain. In such settings, we use the notion of the *transform point spread function* (TPSF). Let Ψ denote the sparsifying transform, and then define $\text{TPSF}(i, j) = (\Psi^* \mathcal{F}_S^* \mathcal{F}_S \Psi)(i, j)$. With this notation, coherence is formally measured by $\max_{i \neq j} |\text{TPSF}(i, j)|$, the maximum size of any off-diagonal entry in the TPSF. Small coherence, e.g., incoherence, is desirable. More discussion about the TPSF can be found in [18].

INCOHERENT MRI ACQUISITION

We now consider several schemes and their associated coherence properties. In two-dimensional (2-D) Cartesian MRI, complete Cartesian sampling is often implemented as a series of n_{pe} acquisitions (called *phase encodes*) along very simple trajec-

tries: parallel equispaced lines. This scheme yields n_{fe} k -space samples per trajectory (called *frequency encodes*), producing a Cartesian grid of $n_{pe} \times n_{fe}$ samples overall. In 3-D, there is an additional encoding dimension (called *slice encode*) requiring $n_{pe} \times n_{se}$ line acquisitions resulting in $n_{pe} \times n_{se} \times n_{fe}$ grid size. The sampling along a trajectory, e.g., the frequency encodes, are rarely a limiting factor in terms of sampling-rate and in terms of the scan time. The number of acquisition lines, e.g., the phase and slice encodes, is limiting.

This immediately suggests speeding up a scan by simply dropping entire lines from an existing complete grid. This is indeed practical: one has complete freedom in choosing the lines to acquire, and the number of lines is what determines the overall scan time, so scan time reduction is exactly proportional to the degree of undersampling. In fact, implementation of this scheme requires only minor modifications to existing pulse sequences—simply skip certain acquisitions. Since most pulse sequences in clinical use are Cartesian, it is very convenient to implement a CS acquisition this way.

Undersampling parallel lines suffers a drawback: the achievable coherence is significantly worse than with truly random k -space sampling. In 2-D imaging, only one dimension is undersampled, so we are only exploiting 1-D sparsity. This reduced incoherence is visible in Figure 6(a). In 3-D Cartesian imaging the situation greatly improves [see Figure 6(b)] since two dimensions are undersampled, and 2-D cross sections are significantly more compressible than their 1-D profiles, so the effectiveness of CS is much higher.

Getting completely away from a Cartesian grid allows far greater flexibility in designing sampling trajectories with low coherence. Popular non-Cartesian schemes include sampling along radial lines or spirals. Traditionally, undersampled radial trajectories [2], [3] and variable density spirals [26] have been used to accelerate acquisitions, because the artifacts from linear reconstruction seem benign and incoherent, much like adding noise to the image. From our perspective, we recognize that such artifacts are benign because *the corresponding PSFs are incoherent*. Figure 6(c)–(f) shows the PSF of several such trajectories. These trajectories are excellent candidates for CS: with appropriate nonlinear reconstruction, the seeming noise-like artifacts can be suppressed without degrading image quality.

A dynamic sequence of images is a multidimensional signal with two or three spatial coordinates and time as an additional dimension (See Figure 7(a), top panel). Dynamic MRI data are acquired in the spatial frequency versus time ($k - t$) domain. Instead of sampling the $k - t$ domain on a regular set of congruent lines (Figure 7(c) top), randomly ordering the lines (Figure 7(c) bottom) randomly samples the $k - t$ space [10], [19] and is incoherent with respect to the spatial versus temporal frequency $x - f$ domain. So random ordering of lines is an effective and inexpensive way to incoherently sample dynamic data. Of

course, the same ideas of random ordering apply to non-Cartesian sampling such as radial lines and spirals, improving incoherence and better exploiting the hardware.

IMAGE RECONSTRUCTION

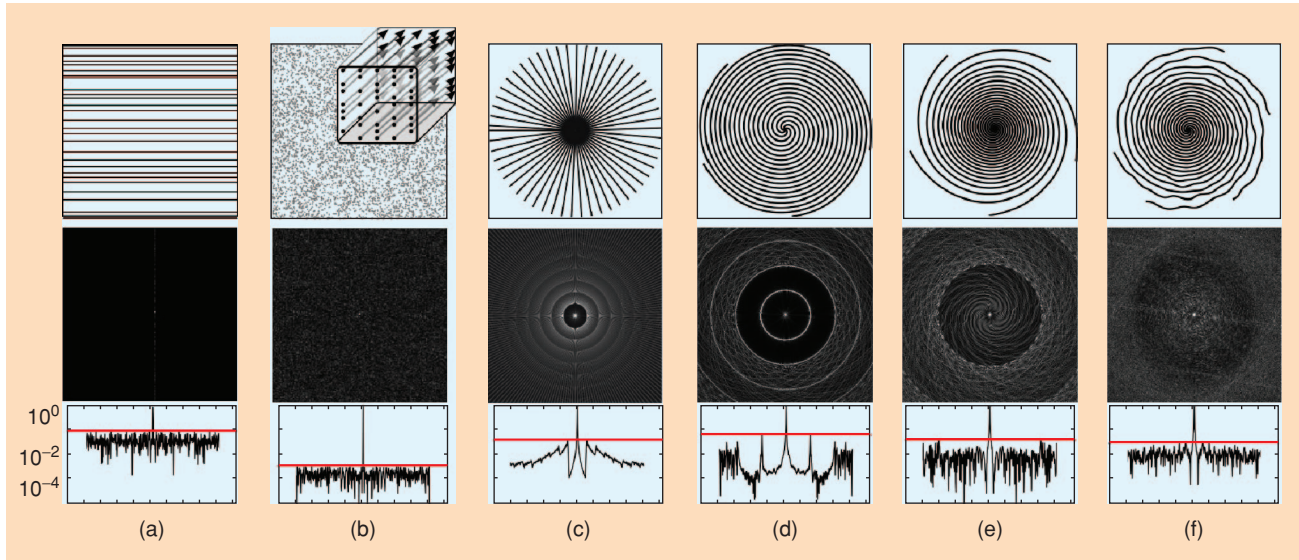
We now briefly describe a useful formal approach for reconstruction. Represent the reconstructed image by a complex vector m , let Ψ denote the linear operator that transforms from pixel representation into the chosen representation. Let \mathcal{F}_S denote the undersampled Fourier transform, corresponding to one of the k -space undersampling schemes discussed earlier. Our reconstructions are obtained by solving the following constrained optimization problem:

$$\begin{aligned} & \text{minimize} && \|\Psi m\|_1 \\ & \text{s.t.} && \|\mathcal{F}_S m - y\|_2 < \epsilon, \end{aligned}$$

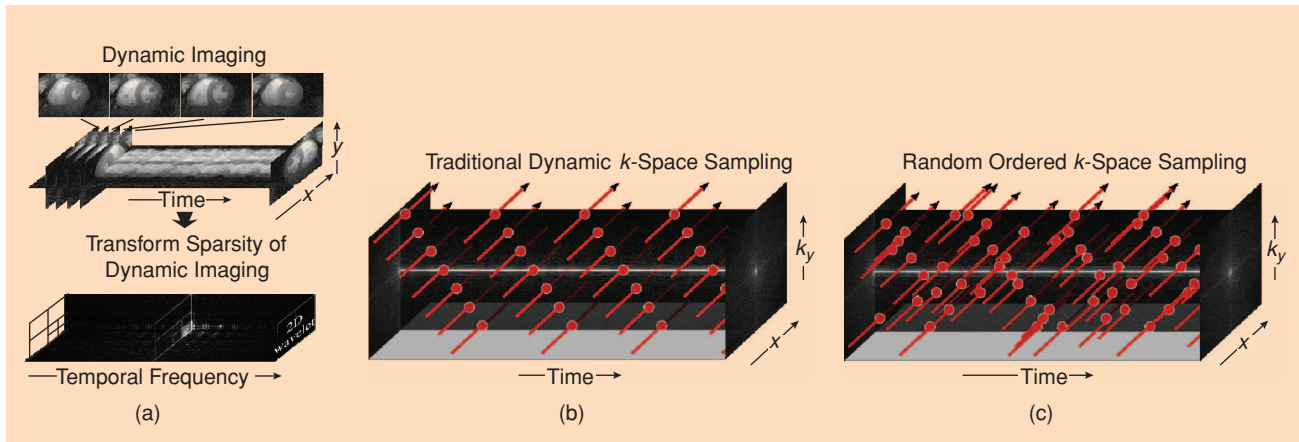
where y is the measured k -space data from the MRI scanner and ϵ controls the fidelity of the reconstruction to the measured data. The threshold parameter ϵ is roughly the expected noise level. Here the ℓ_1 norm $\|x\|_1 = \sum_i |x_i|$.

Minimizing the ℓ_1 norm of $\|\Psi m\|_1$ promotes sparsity [15]–[17]. The constraint $\|\mathcal{F}_S m - y\|_2 < \epsilon$ enforces data consistency. In words, among all solutions that are consistent with the acquired data, we want to find a solution that is compressible by the transform Ψ . It is worth mentioning that when finite-differencing is used as the sparsifying transform, the objective becomes the well-known total variation (TV) penalty [27].

The reader may well ask how such formal optimization-based reconstructions relate to the informal idea of successive interference cancellation. In fact, iterative algorithms for solving such formal optimization problems in effect perform



[FIG6] PSFs of various sampling trajectories: (a) random lines in 2-D, (b) random points in 2-D or cross section of random lines in 3-D, (c) radial, (d) uniform spirals, (e) variable density spirals, and (f) variable density perturbed spirals. The height of the red lines measures coherence.



[FIG7] (a) Dynamic MRI is a multidimensional signal with two or three spatial coordinates and time as an additional dimension. Dynamic images have a sparse representation in an appropriate transform domain. (b) Traditional $k - t$ sequential sampling. (c) Random ordering is an efficient way to incoherently sample the $k - t$ space.

thresholding and interference cancellation at each iteration, so there is a close connection between our exposition and more formal approaches [25], [27], [29].

APPLICATIONS OF COMPRESSED SENSING TO MRI

We now describe four potential applications of CS in MRI. The three requirements for successful CS come together differently in different applications. Of particular interest is the way in which different applications face different constraints, imposed by MRI scanning hardware or by patient considerations, and how the inherent freedom of CS to choose sampling trajectories and sparsifying transforms plays a key role in matching the constraints.

THE CONCEPTS AND APPROACHES WE DISCUSS HERE MAY POTENTIALLY ALLOW ENTIRELY NEW APPLICATIONS OF MRI, ONES CURRENTLY THOUGHT TO BE INTRACTABLE.

RAPID 3-D ANGIOGRAPHY

Angiography is important for diagnosis of vascular disease. Often, a contrast agent is injected, significantly increasing the blood signal compared to the background tissue. In angiography, important diagnostic information is contained in the dynamics of the contrast agent bolus. Capturing the dynamics requires high spatial and temporal resolution of a large FOV, obviously a very difficult task. Today MR angiography scans are often undersampled [3], [11], obtaining improved spatial resolution and temporal resolution at the expense of undersampling artifacts.

CS is particularly suitable for angiography. As shown in Figure 3, angiograms are inherently sparse in the pixel representation and by spatial finite differencing. The need for high temporal and spatial resolution strongly encourages undersampling. CS improves current strategies by significantly reducing the artifacts that result from undersampling.

In this example, we apply CS to 3-D Cartesian contrast-enhanced angiography, which is the most common scheme in clinical practice. Figure 8 illustrates the collection scheme, acquiring equispaced parallel lines in k -space.

Choosing a pseudorandom subset with variable k -space density of 10% of those lines combines undersampling with low coherence. Figure 8 shows a maximum intensity projection (MIP) through the 3-D volume of several reconstructions. CS is able to significantly accelerate MR angiography, enabling better temporal resolution or alternatively improving the resolution of current imagery without compromising scan time.


The nonlinear reconstruction in CS avoids most of the artifacts that appear in linear reconstruction from undersampled data.

WHOLE-HEART CORONARY IMAGING

X-ray coronary angiography is the gold standard for evaluating coronary artery disease, but it is invasive. Multislice X-ray CT is a noninvasive alternative but requires high doses of ionizing radiation. MRI is emerging as a noninvasive, nonionizing alternative.

Multislice X-ray CT is a noninvasive alternative but requires high doses of ionizing radiation. MRI is emerging as a noninvasive, nonionizing alternative.

Coronary arteries are constantly in motion, making high-resolution imaging a challenging task. The effects of heart motion can be minimized by synchronizing acquisitions to the cardiac cycle. The effect of breathing can be minimized by tracking and compensating for respiratory motion or by simply imaging during a short breath-held interval. However, breath-held cardiac-triggered approaches face strict timing constraints and very short imaging windows. The number of acquisitions is limited to the number of cardiac cycles in the breath-hold period. The number of heart-beats per period is itself limited—sick patients cannot be expected to hold their breath for long! Also, each acquisition must be very short to avoid motion blurring. In addition, many slices must be collected to cover the entire heart. These constraints on breath-held cardiac triggered acquisitions traditionally resulted in limited spatial resolution with partial coverage of the heart. Compressed sensing can accelerate data acquisition, allowing the entire heart to be imaged in a single held breath [30].



[FIG8] 3-D Contrast enhanced angiography. Right: Even with 10-fold undersampling CS can recover most blood vessel information revealed by Nyquist sampling; there is significant artifact reduction compared to linear reconstruction; and a significant resolution improvement compared to a low-resolution centric k -space acquisition. Left: The 3-D Cartesian random undersampling configuration.

Figure 9 shows a diagram of the multislice acquisition. To meet the strict timing requirements, an efficient spiral k -space trajectory is used. For each cardiac trigger, a single spiral in k -space is acquired for each slice. The timing limitations still require two-fold k -space undersampling. Therefore we used variable density spirals [26] which have an incoherent PSF [Figure 6(e)], in which linear gridding reconstruction [5] produces artifacts that appear as added noise.

Coronary images are generally piecewise smooth and are sparsified well by finite-differences. CS reconstruction suppresses undersampling-induced interference without degrading the image quality. Figure 9 shows a comparison of the linear direct gridding reconstruction and CS, on the right coronary artery reformatted from a single breath-hold whole-heart acquisition. The linear gridding reconstruction suffers from apparent noise artifacts actually caused by undersampling. The CS reconstruction suppresses those artifacts, without impairing the image quality.

IN MRI, WE LOOK AT A SPECIAL CASE OF CS WHERE THE SAMPLED LINEAR COMBINATIONS ARE SIMPLY INDIVIDUAL FOURIER COEFFICIENTS.

BRAIN IMAGING

Brain scans are the most common clinical application of MRI. Most brain scans use 2-D Cartesian multislice acquisitions. “The Sparsity/Compressibility of MR Images” showed that brain images exhibit transform sparsity in the wavelet domain. The ideas of CS promise to reduce collection time while improving the resolution of current imagery.

We tested the application of CS to brain imaging by acquiring a full Nyquist-sampled data set which we then retrospectively undersampled. The 2-D Cartesian multislice sampling

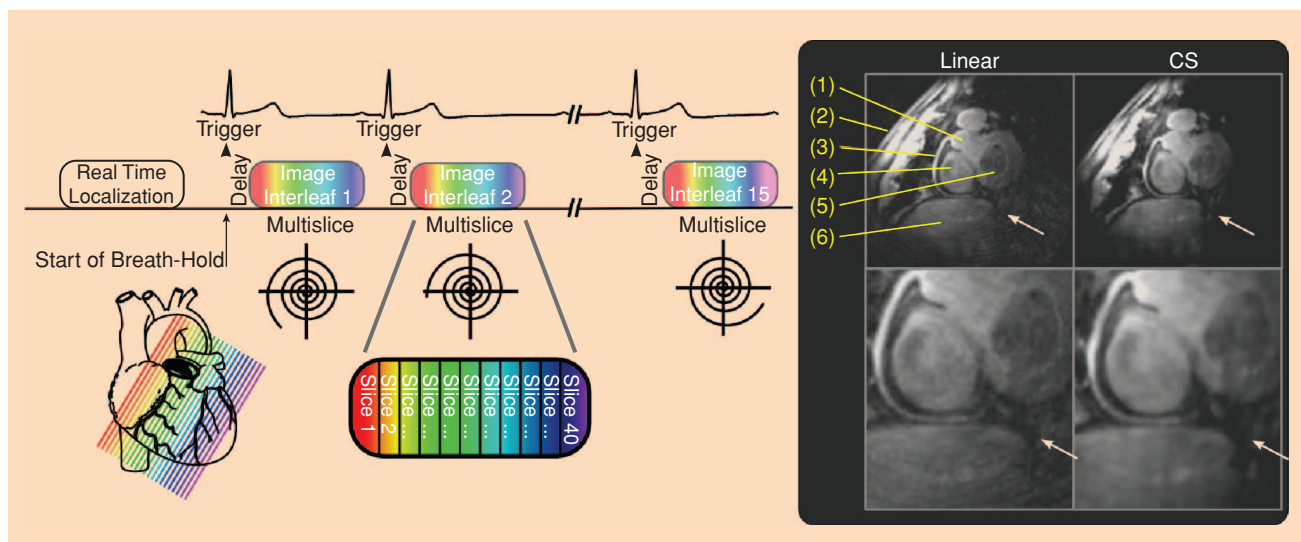
trajectories are illustrated in Figure 10. For each slice we selected a different random subset of 80 trajectories from 192 possible trajectories, a speedup factor 2.4. Undersampling each slice differently reduces coherence compared to sampling the same way in all slices [18].

Figure 10 also shows the experimental results. An axial slice of the multislice CS reconstruction is compared to full Nyquist sampling, linear reconstruction from the undersampled data, and linear reconstruction from a low resolution (LR) acquisition taking the same amount of scan time. CS exhibits both significant resolution improvement over LR at the same scan time, and significant suppression of the aliasing artifacts compared to the linear reconstruction with the same undersampling.


K-T SPARSE: APPLICATION TO DYNAMIC HEART IMAGING

Dynamic imaging of time-varying objects is challenging because of the spatial and temporal sampling requirements of the Nyquist criterion. Often temporal resolution is traded off against spatial resolution. Artifacts appear in the traditional linear reconstruction when the Nyquist criterion is violated.

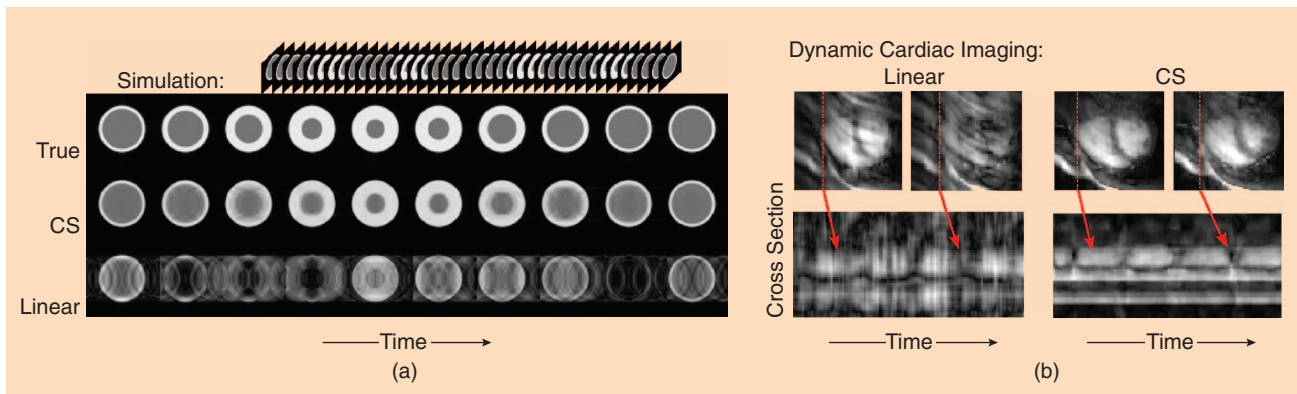
Now consider a special case: dynamic imaging of time-varying objects undergoing quasi-periodic changes. We focus here on heart imaging. Since heart motion is quasi-periodic, the time series of intensity in a single voxel is sparse in the temporal frequency domain (see Figure 3). At the same time, a single frame of the heart “movie” is sparse in the wavelet domain. A simple transform can exploit both effects: apply a spatial wavelet transform followed by a temporal Fourier transform [see Figure 7(a)].



[FIG9] Single breath-hold whole-heart coronary artery imaging. (a) the imaging sequence timing diagram. (b) A slice through the volume of the heart showing the right coronary artery (3). The incoherent artifacts of undersampled variable-density spirals (white arrow) appear as noiselike interference in the linear gridding reconstruction (left). These artifacts are suppressed in the CS reconstruction (right) without compromising image quality. The slice shows: (1) aorta, (2) chest wall, (3) right coronary artery, (4) right ventricle, (5) left ventricle, and (6) liver.



[FIG10] CS exhibits better suppression of aliasing artifacts than linear reconstruction from incoherent sampling, improved resolution over a low-resolution acquisition with the same scan time, and a comparable reconstruction quality to a full Nyquist-sampled set.



[FIG11] Dynamic imaging of quasi-periodic change. (a) Phantom experiment showing a reconstruction from 4-fold undersampling. (b) Dynamic acquisition of the heart motion showing a reconstruction from 7-fold undersampling.

Can we exploit the natural sparsity of dynamic sequences to reconstruct a time-varying object sampled at significantly sub-Nyquist rates? Consider the Cartesian sampling scheme that acquires a single line in k -space for each time slice, following an orderly progression through the space of lines as time progresses [see Figure 7(b)]. For our desired FOV and resolution it is impossible to meet the spatial-temporal Nyquist rate using this scheme.

In fact, this approach is particularly inefficient for dynamic imaging with traditional acquisition and reconstruction methods. Instead, we make one change: make the k -space line ordering *random* instead of orderly [10], [19]. The random ordering comes much closer to randomly sampling $k-t$ space [see Figure 7(c)] and the sampling operator becomes much less coherent with the sparsifying transform.

Figure 11 shows results from two experiments. The first result used synthetic data: a motion phantom, periodically changing in a cartoon of heart motion. The figure depicts an image sequence reconstructed from a sampling rate four times slower than the Nyquist rate, using randomly ordered acquisition and nonlinear reconstruction. The second result involved dynamic real-time acquisition of heart motion. The

given FOV (16 cm), resolution (2.5 mm), and repetition time (4.4 ms) allows a Nyquist rate of 3.6 frames/s. This leads to temporal blurring and artifacts in the traditionally reconstructed image. By instead using random ordering and CS reconstruction we were able to recover the dynamic sequence at the much higher rate of 25 frames/s with significantly reduced image artifact

CONCLUSIONS

We presented four applications where CS improves on current imaging techniques. The concepts and approaches we discussed may potentially allow entirely new applications of MRI—ones currently thought to be intractable.

CS-MRI is still in its infancy. Many crucial issues remain unsettled. These include: optimizing sampling trajectories, developing improved sparse transforms that are incoherent to the sampling operator, studying reconstruction quality in terms of clinical significance, and improving the speed of reconstruction algorithms. The signal processing community has a major opportunity here. There are fascinating theoretical and practical research problems, promising substantial payoffs in improved medical care.

ACKNOWLEDGMENTS

This project was supported by NIH grants P41 RR09784, R01 HL074332, and R01 HL075803 and GE Healthcare. The authors would like to thank Walter Block for his help and Peder Larson, Daeho Lee, Seung-Jean Kim, and Yonit Lustig for their help with the manuscript.

AUTHORS

Michael Lustig (mlustig@stanford.edu) received his B.Sc. in 2001 from Technion-Israel Institute of Technology, Haifa, Israel. In 2008, he received his Ph.D. from Stanford University, where he was working on the application of compressed sensing to rapid MRI. His current research interests include medical imaging reconstruction techniques, MR pulse sequence design, convex optimization, and inverse problems.

David L. Donoho (donoho@stanford.edu) received his A.B. in statistics from Princeton University (Summa Cum Laude), and his Ph.D. in statistics from Harvard University. He has been a Professor of Statistics at both University of California, Berkeley and at Stanford University, and has served as Visiting Professor at the University of Paris, Sackler Visiting Professor at Tel Aviv University, Distinguished Professor at the National University of Singapore, and Kloosterman Professor at the University of Leiden. He is a member of the U.S. National Academy of Science and the American Academy of Arts and Sciences.

Juan M. Santos (jmsantos@mrsrl.stanford.edu) received a B.S. in industrial engineering in 1996 and an M.S. in 1998 from Universidad Catolica de Chile. He received his Ph.D. in electrical engineering from Stanford University in 2006. He has been a research associate in the Stanford Magnetic Resonance System Research Laboratory and is one of the founders of HeartVista, Inc., an NIH SBIR funded startup dedicated to develop cardiovascular MRI applications. His research interests include rapid imaging and reconstruction, real-time MRI, and cardiovascular and interventional applications.

John M. Pauly (pauly@stanford.edu) is an Associate Professor of Electrical Engineering at Stanford University. His main research interests are in MRI, and the use of MRI for guiding minimally invasive interventional procedures. He teaches classes in image reconstruction for medical imaging, as well as RF pulse design for MRI at Stanford. He holds 42 U.S. patents and has authored and co-authored 110 journal articles. He is a Member of the IEEE and is Associate Editor of *IEEE Transactions on Medical Imaging*.

REFERENCES

[1] G. Wright, "Magnetic resonance imaging," *IEEE Signal Processing Mag.*, vol. 14, no. 1, pp. 56–66, Jan. 1997.

[2] K. Scheffler and J. Hennig, "Reduced circular field-of-view imaging," *Magn. Reson. Med.*, vol. 40, no. 3, pp. 474–480, 1998.

[3] D.C. Peters, F.R. Korosec, T.M. Grist, W.F. Block, J.E. Holden, K.K. Vigen, and C.A. Mistretta, "Undersampled projection reconstruction applied to MR angiography," *Magn. Reson. Med.*, vol. 43, no. 1, pp. 91–101, 2000.

[4] C.H. Meyer, B.S. Hu, D.G. Nishimura, and A. Macovski, "Fast spiral coronary artery imaging," *Magn. Reson. Med.*, vol. 28, no. 2, pp. 202–213, 1992.

[5] J.I. Jackson, C.H. Meyer, D.G. Nishimura, and A. Macovski, "Selection of a convolution function for Fourier inversion using gridding," *IEEE Trans. Med. Imag.*, vol. 10, no. 3, pp. 473–478, 1991.

[6] D.K. Sodickson and W.J. Manning, "Simultaneous acquisition of spatial harmonics (SMASH): Fast imaging with radiofrequency coil arrays," *Magn. Reson. Med.*, vol. 38, no. 4, pp. 591–603, 1997.

[7] K.P. Pruessmann, M. Weiger, M.B. Scheidegger, and P. Boesiger, "SENSE: Sensitivity encoding for fast MRI," *Magn. Reson. Med.*, vol. 42, no. 5, pp. 952–962, 1999.

[8] N.P.B. Madore, G.H. Glover, "Unaliasing by fourier-encoding the overlaps using the temporal dimension (UNFOLD), applied to cardiac imaging and fMRI," *Magn. Reson. Med.*, vol. 42, no. 5, pp. 813–828, 1999.

[9] N.P. Willis and Y. Bresler, "Optimal scan for time-varying tomography. I. Theoretical analysis and fundamental limitations," *IEEE Trans. Image Processing*, vol. 4, no. 5, pp. 642–653, 1995.

[10] T. Parrish and X. Hu, "Continuous update with random encoding (CURE): A new strategy for dynamic imaging," *Magn. Reson. Med.*, vol. 33, no. 3, pp. 326–336, Mar. 1995.

[11] F.R. Korosec, R. Frayne, T.M. Grist, and C.A. Mistretta, "Time-resolved contrast-enhanced 3D MR angiography," *Magn. Reson. Med.*, vol. 36, no. 3, pp. 345–351, Sept. 1996.

[12] J.M. Hanson, Z.P. Liang, R.L. Magin, J.L. Duerk, and P.C. Lauterbur, "A comparison of RIGR and SVD dynamic imaging methods," *Magn. Reson. Med.*, vol. 38, no. 1, pp. 161–167, July 1997.

[13] J. Tsao, P. Boesiger, and K.P. Pruessmann, "k-t BLAST and k-t SENSE: Dynamic MRI with high frame rate exploiting spatiotemporal correlations," *Magn. Reson. Med.*, vol. 50, no. 5, pp. 1031–1042, 2003.

[14] C.A. Mistretta, O. Wieben, J. Velikina, W. Block, J. Perry, Y. Wu, K. Johnson, and Y. Wu, "Highly constrained backprojection for time-resolved MRI," *Magn. Reson. Med.*, vol. 55, no. 1, pp. 30–40, Jan. 2006.

[15] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[16] D. Donoho, "Compressed sensing," *IEEE Trans. Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.

[17] E. Candès and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.

[18] M. Lustig, D.L. Donoho, and J.M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, 2007.

[19] M. Lustig, J.M. Santos, D.L. Donoho, and J.M. Pauly, "k-t Sparse: High frame rate dynamic MRI exploiting spatio-temporal sparsity," in *Proc. 13th Annu. Meeting ISMRM*, Seattle, WA, 2006, p. 2420.

[20] J.C. Ye, S. Tak, Y. Han, and H.W. Park, "Projection reconstruction MR imaging using FOCUS," *Magn. Reson. Med.*, vol. 57, no. 4, pp. 764–775, Apr. 2007.

[21] K.T. Block, M. Uecker, and J. Frahm, "Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint," *Magn. Reson. Med.*, vol. 57, no. 6, pp. 1086–1098, June 2007.

[22] T.-C. Chang, L. He, and T. Fang, "MR image reconstruction from sparse radial samples using Bregman iteration," in *Proc. 13th Annu. Meeting ISMRM*, Seattle, WA, 2006, p. 696.

[23] F. Wajer, "Non-Cartesian MRI scan time reduction through sparse sampling," Ph.D. dissertation, Dept. Appl. Phys., Delft University of Technology, The Netherlands, 2001.

[24] D.S. Taubman and M.W. Marcellin, *JPEG 2000: Image Compression Fundamentals, Standards and Practice* (International Series in Engineering and Computer Science), Norwell, MA: Kluwer, 2002.

[25] D. Donoho, Y. Tsaig, I. Drori, and J.-L. Starck, "Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit," Tech. Report 06-02, Dept. of Statistics, Stanford Univ., 2006.

[26] C.-M. Tsai and D. Nishimura, "Reduced aliasing artifacts using variable-density k-space sampling trajectories," *Magn. Reson. Med.*, vol. 43, no. 3, pp. 452–458, 2000.

[27] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation noise removal algorithm," *Phys. D*, vol. 60, pp. 259–268, Nov. 1992.

[28] E. Candès and J. Romberg, "Signal recovery from random projections," in *Proc. SPIE Comp. Imaging III*, San Jose, CA, 2005, p. 5674.

[29] M. Elad, B. Matalon, and M. Zibulevsky, "Coordinate and subspace optimization methods for linear least squares with non-quadratic regularization," *J. Appl. Comput. Harmon. Anal.*, vol. 23, pp. 346–367, Nov. 2007.

[30] J.M. Santos, C.H. Cunningham, M. Lustig, B.A. Hargreaves, B.S. Hu, D.G. Nishimura, and J.M. Pauly, "Single breath-hold whole-heart MRA using variable-density spirals at 3T," *Magn. Reson. Med.*, vol. 55, no. 2, pp. 371–379, 2006. 