Application of “Compressed Sensing” for Rapid MR Imaging

Michael Lustig[1], Juan M. Santos[1], Jin-Hyung Lee[1], David L. Donoho[2] and John M. Pauly[1]
[1] Electrical Engineering Department, Stanford University, CA USA [2] Statistics Department, Stanford University, CA, USA

Introduction

In a typical magnetic resonance imaging (MRI) experiment, samples are collected in the so-called k-space or frequency domain of the image. The number of samples needed for reconstruction at a given resolution and field of view is normally set by the Nyquist criteria and occupies a certain amount of bits in memory. However, it is commonly known that if we take the reconstructed image and compress it either losslessly or visually losslessly using a compression algorithm, we can represent the exact same image with fewer bits. Recent theoretical results in sparse signal recovery show that if the underlying image is compressible it can be recovered from randomly undersampled frequency data, an idea also known as “compressed sensing” [1,2]. Random sampling per-se is impractical and inefficient for MRI hardware. In this work, we develop practical under sampling schemes for MR imaging by adapting common acquisition methods. Specifically we show examples for Cartesian and spiral imaging. We introduce randomness by perturbing the k-space trajectories. We reconstruct by minimizing the L_1 norm of a transformed image subject to data fidelity constraints. Simulations and experimental results show good reconstruction particularly from heavily undersampled k-space data where conventional methods fail. We further demonstrate the application of this scheme for a single breath hold whole heart coronary angiography [8], where the scan time speedup has tremendous clinical significance.

Theory

An MRI scan consists of series of repeated short experiments. In each experiment, different part of the Fourier-space (k-space) information of the image is collected. The experiments continue until a full data set is gathered. Because of physical MR hardware limitations, the k-space information in each repetition is collected along lines or smooth curves. Typical repetition numbers vary between a few to hundreds (in 2D imaging) and thousands (3D imaging) of repetitions. The numbers vary depending on the application, resolution, field of view, SNR and the trajectory in which the samples are collected. Typical repetition times (also known as TR) range from a few ms to a few seconds, again depending on the application. In general, the total scan time is highly dependent on the number of samples needed for reconstruction. In practice, it is more dependent on the number of repetitions rather than the total number of samples. Therefore, by reducing the total number of repetitions, the gain in scan time reduction is more significant than reducing the number of samples within a single repetition.

The most common way to sample k-space is along lines on a Cartesian grid. This form of sampling is used routinely mainly because of its simplicity and robustness to hardware and physical imperfections. In this type of sampling, one or more lines are collected in each TR. Spiral imaging is another form of k-space sampling. It is known for its hardware efficiency, fast imaging and robustness to flow. Spirals are mainly used for real-time and fast imaging applications. In this type of imaging samples along a spiral trajectory are collected every TR. Each spiral is rotated by an angle such that N spirals will sufficiently cover k-space. Radial sampling is another common way of covering k-space, however, it can be thought of as a private case of variable density spirals.

The goal of undersampled reconstruction is to reconstruct an image from incomplete Fourier data -- a highly under-determined problem. Medical images are often compressible by a linear transform (such as finite differences, wavelets, etc.), where the number of coefficients needed to describe the image accurately is significantly smaller than the number of pixels in the image. We exploit compressibility by constraining our reconstruction to have a sparse representation and be consistent with the measured k-space data. Surprisingly, if the underlying true object has a sparse representation (see [1,2,3] for details) and the sampling in frequency is uniformly and randomly distributed, we can recover the signal accurately by solving the following constrained optimization problem,

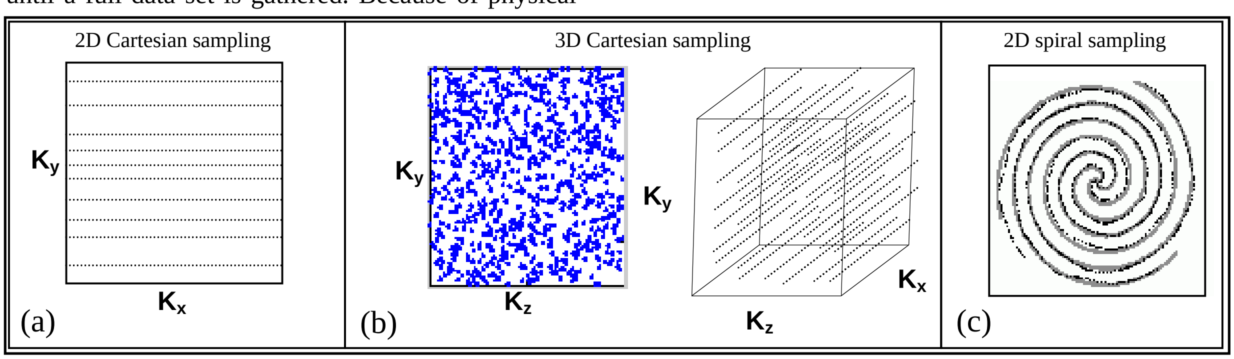

Figure 1: (a) Randomized 2D Cartesian sampling. Undersampling is done by randomly removing lines and randomly perturbing the remaining ones about their location. (b) Randomized 3D Cartesian sampling. Undersampling is done by randomly removing lines and randomly perturbing the remaining ones in the k_y-k_z plane. (c) Spiral and perturbed spiral sampling. Randomness is introduced by deviating from the analytic spiral and slightly perturbing the angles between the interleaves

$$\text{minimize } ||\Psi(m)||_1 \quad (1)$$
$$\text{s.t } ||Fm - y||_2 < \varepsilon$$

Here, m is the image, Ψ transforms the image into a sparse representation, F is an undersampled Fourier matrix, y is the measured k-space data and ε controls fidelity of the reconstruction to the measured data. ε is usually set to the noise level. The objective enforces sparsity whereas the constraint enforces data consistency. Eq. 1 is a convex quadratic program (QP)[4] for real valued images, and a second order cone program (SOCP) for complex images, each of which have efficient solvers. Note that the Fourier operator is not limited to grid sampling and is in fact a non-uniform Fourier transform operator.

Methods

We propose two different sparsifying transforms, the wavelet transform and finite differences -- both widely used in image processing [7]. For finite differences, the objective becomes the total variation TV = $\Sigma_x \Sigma_y \, |\nabla m(x,y)|$ where $\nabla m(x,y)$ is the spatial gradient of the image, computed by finite differences.

Random sampling, as advocated in [1,2,3] is highly inefficient in MR because of hardware limitations. Instead of randomness, we would like to create irregularity in the sampling pattern by perturbing commonly used trajectories. Doing so, we maintain the efficiency of the scan and also helping the reconstruction by using this "pseudo-random" or irregular type of sampling. In this work we will focus on adapting Cartesian and spiral sampling, however, other forms of sampling can be adapted in the same manner.

In Cartesian imaging, it can be shown that by regularly under sampling, ie. decimating every other line or column, the reconstruction is bound to fail [1]. Instead, we under sample in an irregular fashion by randomly removing lines. We add more randomness by slightly shifting the remaining lines about their location randomly. Fig. 1a and 1b show the under sampling pattern for 2D and 3D imaging respectively. In the 2D case, randomness is introduced only in the y-axis whereas in the 3D case, randomness is introduced in the y-z plane.

Spiral trajectories, uniform and variable density, are good candidates for approximating random sampling. They span k-space uniformly but on the other hand they are far from being as regular as a Cartesian grid. Furthermore spiral imaging is fast and time-efficient. To introduce more randomness we perturbed the individual spiral trajectories, slightly deviating from the deterministic spiral along each interleave; the interleave angles are also perturbed by a small random angle. Fig. 1c shows the sampling pattern of a perturbed spiral trajectory compared to a deterministic spiral.

Phantom and in vivo results

To validate our approach we considered a 34 interleave perturbed spiral trajectory, designed for a 16 cm FOV 1 mm resolution. We undersampled by 45% by acquiring data only on a subset of 19 out of the 34 interleaves using a GRE sequence (TE=1.3ms, TR=8.24, RO=3ms, α=30°, slice= 4mm). The experiment was conducted on a 1.5T GE Signa scanner with gradients capable of 40mT/m and 150mT/m/ms maximum slew rate. The image was reconstructed by TV reconstruction implemented with finite derivatives, and with L_1 wavelet (Daubechies 4) reconstruction. Results were compared to gridding and minimum-norm reconstructions. Our reconstructions used a non-linear conjugate gradient solver [6] with min-max nuFFT [5,6] as the non-uniform Fourier transform engine.

Fig 2. illustrates the results of four reconstruction algorithms for the phantom experiment. As expected, the gridding and minimum norm reconstructions exhibit severe aliasing artifacts due to undersampling. On the other hand, the L_1 reconstructions removed most aliasing artifacts while preserving resolution. TV penalization performs slightly better than the L_1/wavelet penalty. This difference is attributable to the object being piece wise constant and hence sparser for the finite difference operator than for the wavelet transform. Fine structures that are severely corrupted by aliasing are well recovered by the L_1 reconstructions. Note that the Fourier transform of all the reconstructed images in Fig.2 are the same (up to noise level) at the spiral sample points. The L_1 method was able to recover the information because the correct image representation is sparse, and sparsity is being imposed.

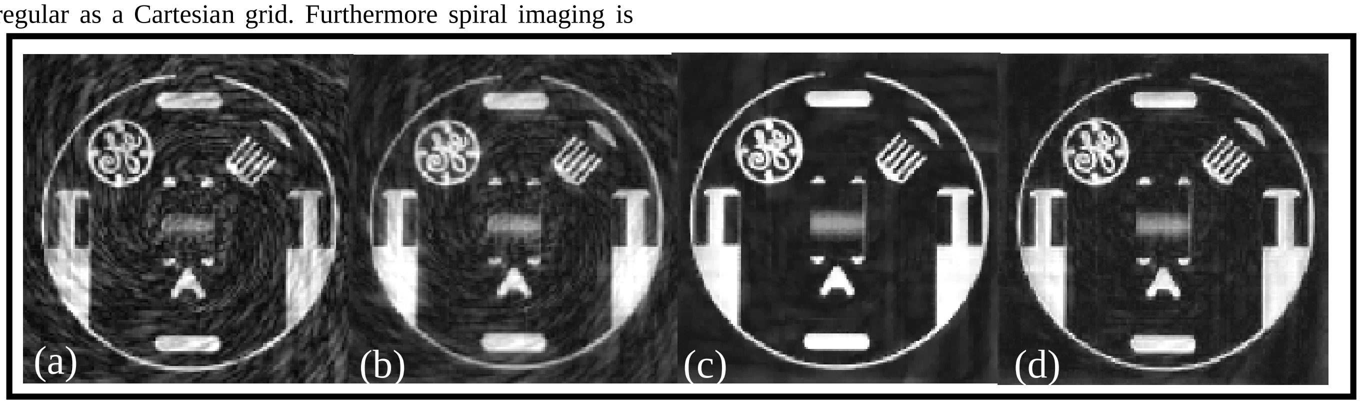

Figure 2: ***Various reconstructions from 55% k-space undersampling. a) gridding b) Minimum norm c) Total variation d) L1 wavelet Note that structures that are severely corrupted by aliasing are recovered by the TV and L1 /wavelet reconstructions.***

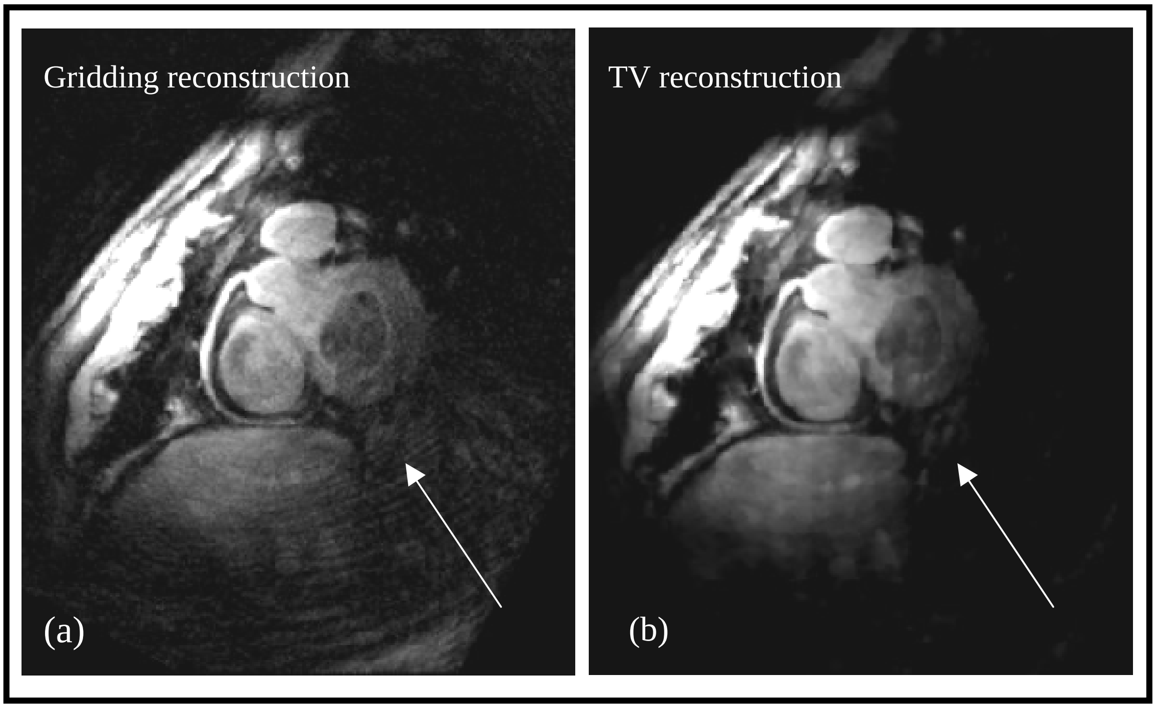

Figure 3: Right coronary artery image acquired using a whole heart, single breath-hold acquisition with variable density spirals. (a) Gridding reconstruction. (b) Total Variation (TV) reconstruction. Note, that the aliasing artifacts due to the undersampling in the gridding reconstruction are removed in the TV reconstruction while maintaining the high resolution and improving the quality of the image.

To further validate our approach we applied our method for whole heart, single breath-hold coronary artery imaging. We designed a 17 interleave variable density spiral with 50% under sampling ratio. The spiral was designed for a nominal FOV of 20 cm and 0.8 mm in plane resolution. The acquisition was done in a multi-slice technique with a total scan time of 17 heartbeats to cover the volume of the heart. Each slice was reconstructed using TV reconstruction. The data was then reformatted to produce a slice showing the right coronary artery. The result was compared to a linear gridding reconstruction of the same data set.

Fig. 3 compares the TV reconstruction of the right coronary artery to the linear gridding reconstruction. The aliasing artifacts from the under sampling in the gridding reconstruction are removed in the TV reconstruction while maintaining the resolution and features, and resulting in a higher quality image.

Conclusion

L_1 –penalized image reconstruction outperforms conventional linear reconstruction, recovering the image even with significant undersampling. The non-linearity of the L_1 norm is the key to the reconstruction; however our method is more computationally intensive than traditional linear methods. In the current, rather inefficient Matlab^{TM} implementation we are able to reconstruct a 256x256 2D image in a matter of several minutes. Our simulations show that using perturbed trajectories offer better reconstruction than just by uniformly undersampling k-space. This type of reconstruction can be used to speed up acquisition whenever there is sparsity to exploit in space or time. Applications such as angiography, time-resolved and contrast enhanced imaging are perfect candidates as such images can be have a very sparse representation, yet the clinical applications are currently significantly limited by imaging times.

References

[1] Candès et al. 2004 “Robust uncertainty principles: Exact Signal Reconstruction from Highly incomplete Frequency Information.” Manuscript: http://www.math.ucla.edu/~tao/preprints/Exact4.pdf

[2] Donoho D.L. 2004, “Compressed Sensing”. Manuscript http://www-stat.stanford.edu/~donoho/Reports/2004/CompressedSensing091604.pdf

[3] Tsaig et al. 2004, “Extensions of Compressed Sensing”, Manuscript http://www-stat.stanford.edu/~donoho/Reports/2004/ExtCS-10-22-04.pdf

[4] Chen et al. SIAM 1999,20:33-61

[5] Fessler et al. IEEE TSP 2003,51:560-57

[6] Bronstein et al. IEEE TMI 2002, 21:1395-1401

[7] Rudin et al. Phys. D 1992, 60:259-268

[8] Santos et al. “Single Breath-Hold Whole-Heart MRA Using Variable-Density Spirals at 3T” 2005, to appear in Mag Med Res