Optimal Sampling in Compressed Sensing

Joyita Dutta

Introduction

Compressed sensing allows us to recover objects reasonably well from highly undersampled data, in spite of violating the Nyquist criterion. In the context of MRI, this technique is appealing since a smaller dataset implies faster imaging. This work seeks to explore the mathematics of compressed sensing and to look for ways to generate optimal sampling patterns. The content of this report includes the following:

- A simplified overview of compressed sensing
- Optimality measures for sampling patterns
- Methodology used for optimally sampling 16x16 images
- Analyses and simulation results based on optimal sampling of 16x16 images

Overview

In order to recover an object $x \in \mathbb{R}^N$ from a measurement set $f \in \mathbb{R}^M$ is where $M << N$, we have to solve an extremely ill-posed problem:

$$\Phi_u x = f \,, \tag{1}$$

where $\Phi_u \in \mathbb{R}^{MxN}$ is the measurement system (the undersampled Fourier basis in the case of MRI). This is an underdetermined system and has infinitely many solutions. So how do we make this problem tractable?

The success of compressed sensing rests on two key tenets. The first key idea of compressed sensing is to add more information to the inverse problem by exploiting the idea that we can find a basis $\Psi \in \mathbb{R}^{NxN}$, in which, the objects we intend to image will have a *sparse* representation. In other words,

$$x = \Psi s \,, \tag{2}$$

where $s \in \mathbb{R}^N$ and has at most K non-zero elements. How this extra information is added will be explained shortly, in the context of the inverse problem formulation.

Before that, we are still to find a way to prevent aliasing artifacts in the reconstructed image due to severe undersampling in the dataset f. This leads us to the second key idea, which is to make sure that the systems Φ_u and Ψ are *incoherent*. This ensures that the aliasing artifacts will appear noise-like in the Ψ domain and are spread out over all the bases. Thus, the K dominant elements of s are less likely to be corrupted by aliasing.

Now, let us look at the formulation of the complete inverse problem:

$$\begin{aligned} &minimize && \left\|\Psi^* x\right\|_1 \\ &s.t. && \left\|\Phi_u x - f\right\|_2 < \varepsilon \end{aligned} \tag{3}$$

The first condition involving the minimization of the L_1 norm of $s = \Psi^* x$ imposes the sparseness condition as illustrated in Figure 1 [1]. The second condition, where ε is a small positive number, ensures that the solution is consistent with the data f.

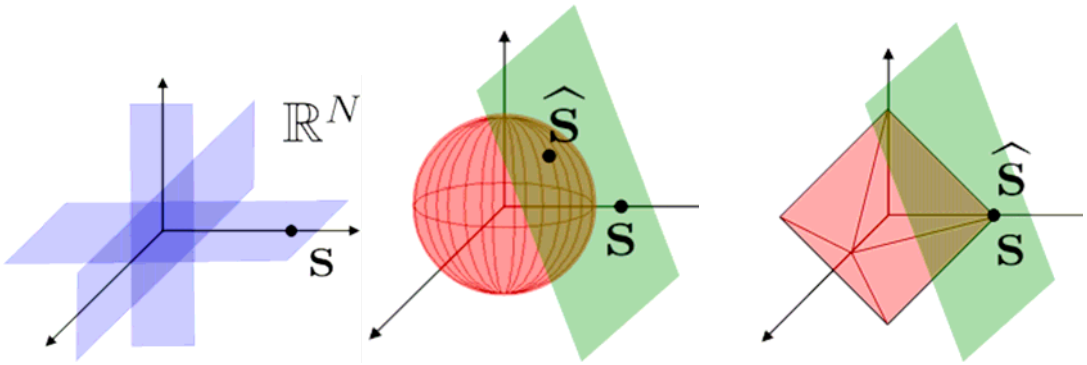

Figure 1: (a) A sparse vector s lies on a K-dimensional hyperplane aligned with the coordinate axes in $\mathbb{R}^N$ and thus close to the axes. (b) L_2 minimization gives us the point where the largest hypersphere touches the nullspace (green hyperplane), (c) L_1 minimization will give us a point near the coordinate axes, which is where the sparse vector *s* is located. [1]

Optimal Sampling

The primary question this work is aiming to tackle is this: How do we choose the best sampling pattern and hence the undersampling measurement operator, Φ_u? The existing literature point at two main approaches:

- *Correlation-based approach*
  Assuming Φ, $\Psi \in \mathbb{R}^{NxN}$ are orthobases, the coherence between the two is given by equation (4) [2,3].

$$\mu(\Phi, \Psi) = \sqrt{N} \max_{1 \le k, j \le N} \left|\left\langle \phi_k, \psi_j \right\rangle\right| \tag{4}$$

  The idea is to choose Φ for a fixed Ψ such that they are maximally incoherent, for reasons mentioned before. It can be shown that $\mu(\Phi, \Psi) \in \left[1, \sqrt{N}\right]$. Also, random matrices are largely incoherent with any given basis Φ. Herein lies the success of random sampling schemes.
  Another way to measure the coherence is to look at the maximum side-lobe-to-peak ratio in the transform point spread function defined in [4] as:

$$TPSF(i,j) = e_j^* \Psi^* \Phi^* \Phi \Psi e_i \ ,$$

(5)

The two measures have similar connotations and similar computational costs (as we shall see later).

- *Combinatorial approach*
  As illustrated in Figure 2, we may look at the measurement, f, as a linear combination of K columns of $\Theta = \Phi_u \Psi$, where $\Theta, \Phi_u \in \mathbb{R}^{MxN}$, $\Psi \in \mathbb{R}^{NxN}$. If we knew beforehand which of the K elements of s are non-zero, we could have constructed a stable, well-conditioned system Θ , by making sure it preserves the lengths of the particular K-sparse vectors, as given by the restricted isometry property [2] in equation (6):

$$\delta_s = \left| 1 - \frac{\|\Theta s\|_2}{\|s\|_2} \right| \to 0+$$

(6)

In real life, we do not know the particular K non-zero locations. Thus we would have to verify equation (6) for every one of the ${}^N C_K$ groups of columns in every possible $\Phi_u \in \mathbb{R}^{MxN}$. This appears to be a daunting task considering the prohibitively expensive computations involved!

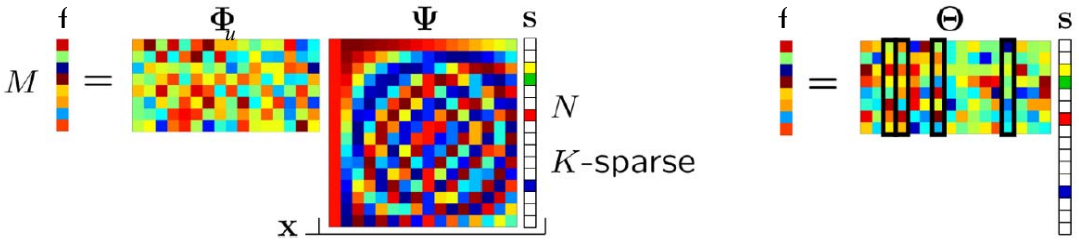

Figure 2: Representation of *f* as a linear combination of *K* columns of $\Theta = \Phi\Psi$, where $\Theta, \Phi \in \mathrm{R}^{MxN}$, $\Psi \in \mathrm{R}^{NxN}$, and *s* is *K*-sparse. [1]

Methodology

Let us now look at the results presented so far in the context of MRI. To simplify computation, let us assume we are trying to recover a small 16x16 image (a total of N = 256 pixels) from undersampled Fourier (k-space) data. Our focus is on a 2DFT Cartesian readout scheme, widely used in MRI, in which undersampling is restricted to the phase-encodes, while the readouts are fully sampled. The 2D Fourier and Wavelet transforms F and S of the image $X \in \mathbb{R}^{nxn}$ (n = 16) may be written as:

$$F = \Phi_{1D} X \Phi_{1D}^*$$
$$S = \Psi_{1D}^* X \Psi_{1D}$$

(7)

Here $\Phi_{1D} \in \mathbb{R}^{nxn}$ and $\Psi_{1D} \in \mathbb{R}^{nxn}$ are the 1D Fourier and Wavelet operators respectively. In order to be able to apply equations 1-6 to our problem, we must first represent the 2D

transforms in the equations (7) as linear transformations of the stacked image vector $x \in \mathbb{R}^N$ (N=256). We use the fact that the separable 2D transform matrix can be written as the Kronecker tensor product of the 1D transforms:

$$\Phi = \Phi_{1D} \otimes \Phi_{1D}$$
$$\Psi = \Psi_{1D} \otimes \Psi_{1D} \quad (8)$$

Once these operators have been constructed, we can formulate our problem in the form described in Figure 2.

We pose our analysis as a series of questions, which will be answered in the next section:

1. For given N and Φ (the full Fourier operator in this case), what Ψ should we chose?
2. What are the computational costs involved in using the two different approaches for finding optimal patterns for a given object size N?
3. For given N, Φ, and Ψ, and for M samples, what Φ_u should be optimal?

Results

The Wavelet Transform

We look for an answer to question 1 assuming N = 256 and Φ is the full (orthonormal) Fourier matrix. The answer will decide two key things: the sparsity, K, of s and the natural coherence $\mu(\Phi, \Psi)$ between the two bases. The first of these is primarily determined by the coarseness level. Figure 3 shows our 16x16 test object, and its level 2 Haar transform. Figure 4 shows plots of the level 2, 3, and 4 Haar coefficients (the level 4 coefficients being the pixel intensities in the object itself). Figure 5 shows the level 2 Coiflet-2 transform coefficients and the image reconstructed from the 15 largest of these indicating the compressibility of our test object.

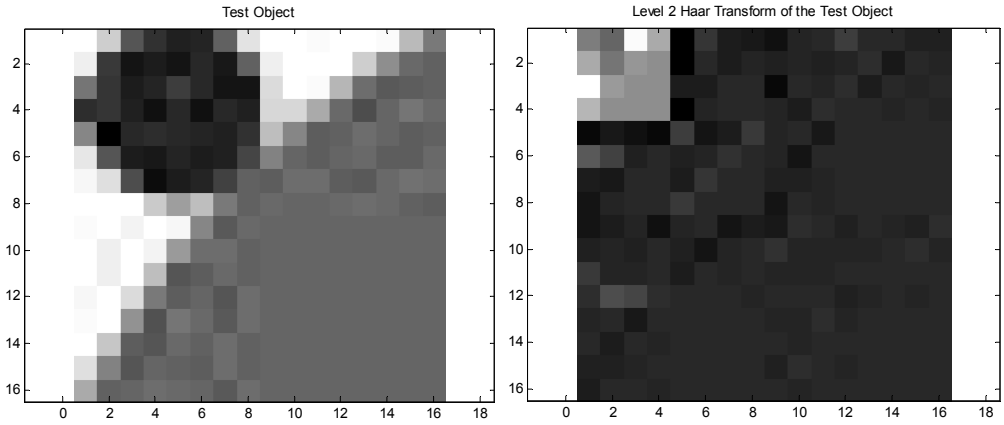

Figure 3: The 16x16 test object (left) and its level 2 Haar transform (right)

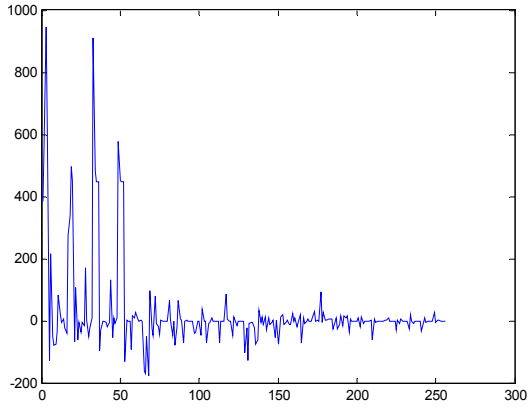

(a)

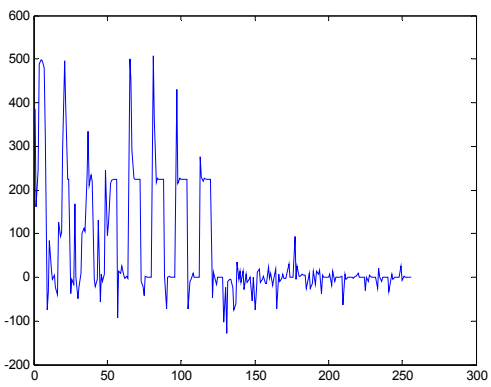

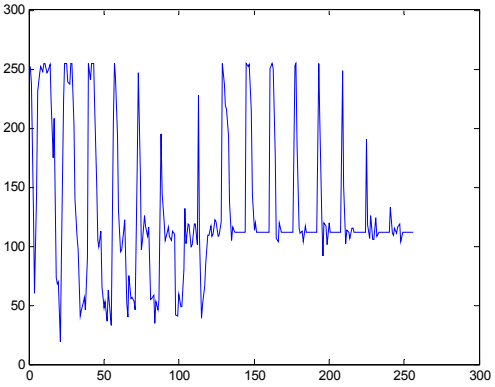

(b) (c)

Figure 4: Haar Coefficients for levels (a) 2, (b) 3, and (c) 4 (actual pixel intensities) respectively

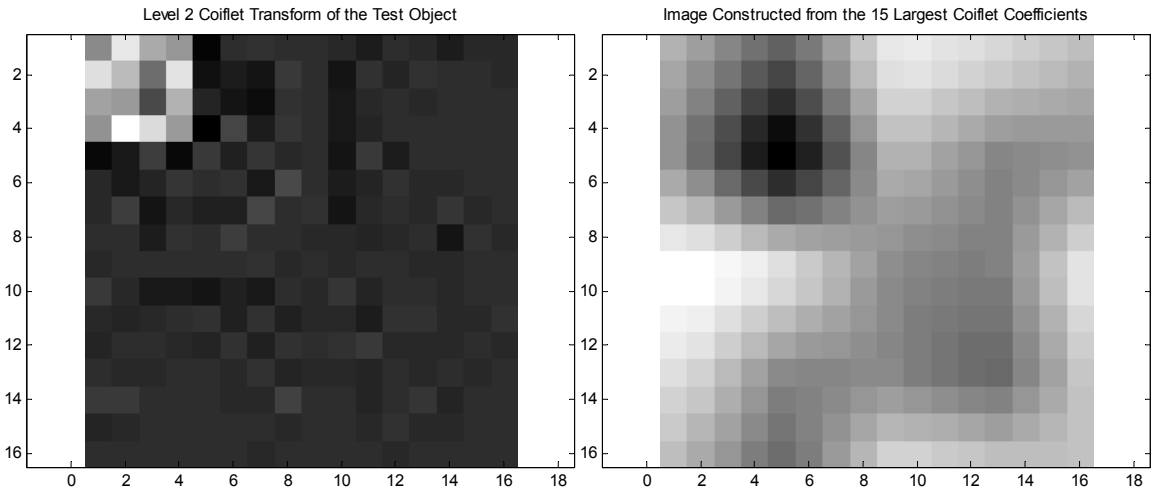

Figure 5: (a) Level 2 transform using Coiflet-2 wavelets (b) Image constructed from the 15 largest coefficients indicating compressibility

The coefficients were computed for different wavelets generated using WaveLab 850, and it was observed that the fraction of nonzero coefficients was about 7% for level 2 and about 20% with level 3. Next we use equation 4 to compute the coherence of different wavelet bases with the Fourier bases. The results are shown in tables I and II for N = 256 and N = 1024 respectively.

Table I

Coherence measures between wavelet bases and the Fourier bases for N = 256

Level	*K*	Haar-2	Coiflet-2	Daubechies-4	Symmlet-4
2	~7%	4	4	4	3.999
3	~20%	2	2	2	1.999
4	-	1	1	1	1.000

Table II

Coherence measures between wavelet bases and the Fourier bases for N = 1024

Level	Haar-2	Coiflet-2	Daubechies-4	Symmlet-4
2	8	8	8	7.999
3	4	4	4	3.999
4	2	2	2	1.999
5	1	1	1	1.000

The analysis gives us similar results for different wavelet types. Also, it indicates a tradeoff between sparsity and incoherence. So we shall be looking at optimal sampling patterns for both levels 2 and 3. Also, we choose the Haar bases for their inherent simplicity and ease of implementation.

Computational Costs

Next, we address question 2 and look at the computational costs involved in using the two different approaches for finding optimal patterns for a given object size N. Table III lists the number of iterations in the combinatorial approach and the correlation-based approaches.

Table III

Number of iterations required by different approaches

Number of ways to pick m samples from n: ${}^{n}C_{m}$	Combinatorial	Coherence	SPR
${}^{16}C_{4}$ = 1820	${}^{256}C_{16} \sim 10^{25}$	65536	65280
${}^{32}C_{8}$ = 10518300	${}^{1024}C_{32} \sim 10^{60}$	1048576	1047552

In order to handle larger images and more phase encodes, we would have to resort to Monte Carlo simulations as was done in [4]. Currently, we are restricting ourselves to a 16x16 image and 4 phase encodes.

Optimal Sampling Patterns

The coherence computations revealed a minimum coherence level of 1. There was a set of patterns corresponding to this coherence level. “Masks” for the corresponding

sampling patterns were generated and the Fourier matrix, Φ, was zero-padded. This zero-padded matrix is then used to generate undersampled k-space data and to reconstruct the object. The Sparse MRI V0.2 set of Matlab functions was used for reconstruction. Figure 6 shows the first test object. Figure 7 shows a set of phase encode levels that lead to minimum coherence and lastly the reconstructed image.

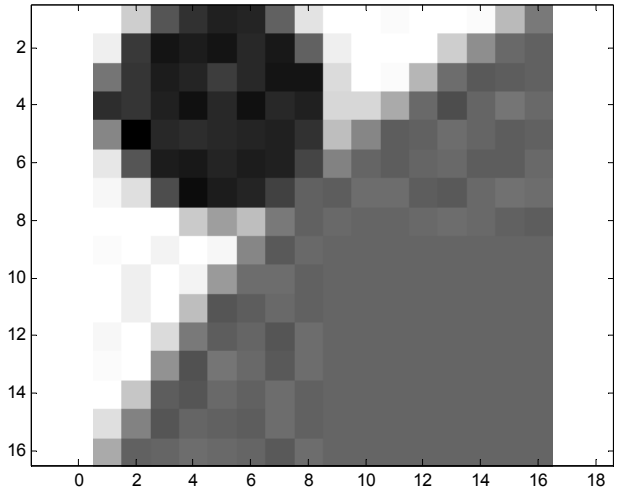

Figure 6: The first test object (same as the one used before for analyzing wavelets)

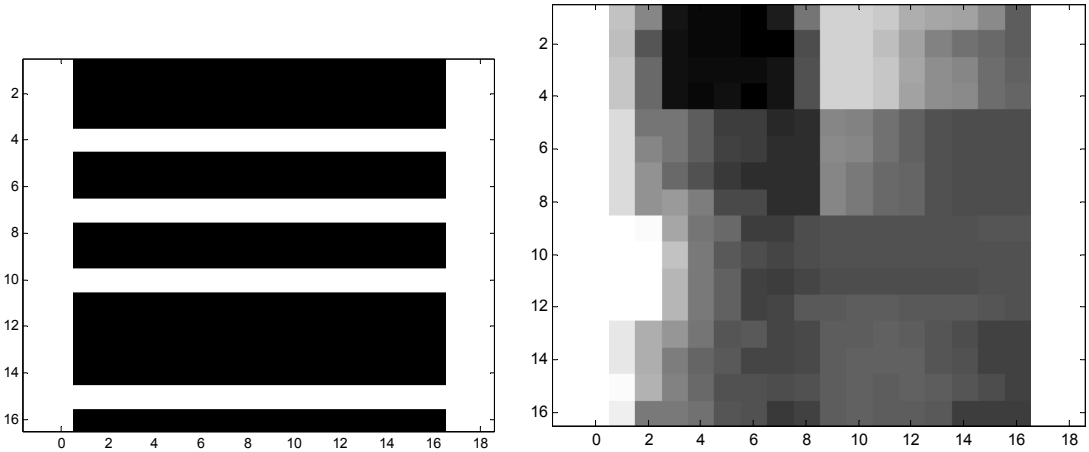

Figure 7: (a) A set of 4 phase encode levels that lead to minimum coherence, and (b) the reconstructed image.

The reconstructed image does not look too good obviously because data around center of k-space is missing. To get around this, we pick a different sampling pattern that gives the same minimum coherence, but includes the central phase encode levels. Figure 8 shows this sampling pattern, the initial guess generated from a backward projection of the zero-padded Fourier matrix on the undersampled data, and lastly the reconstructed image. A second test object is shown in Figure 9. The reconstruction results for this object for the sampling pattern in Figure 8(a) is shown in Figure 10.

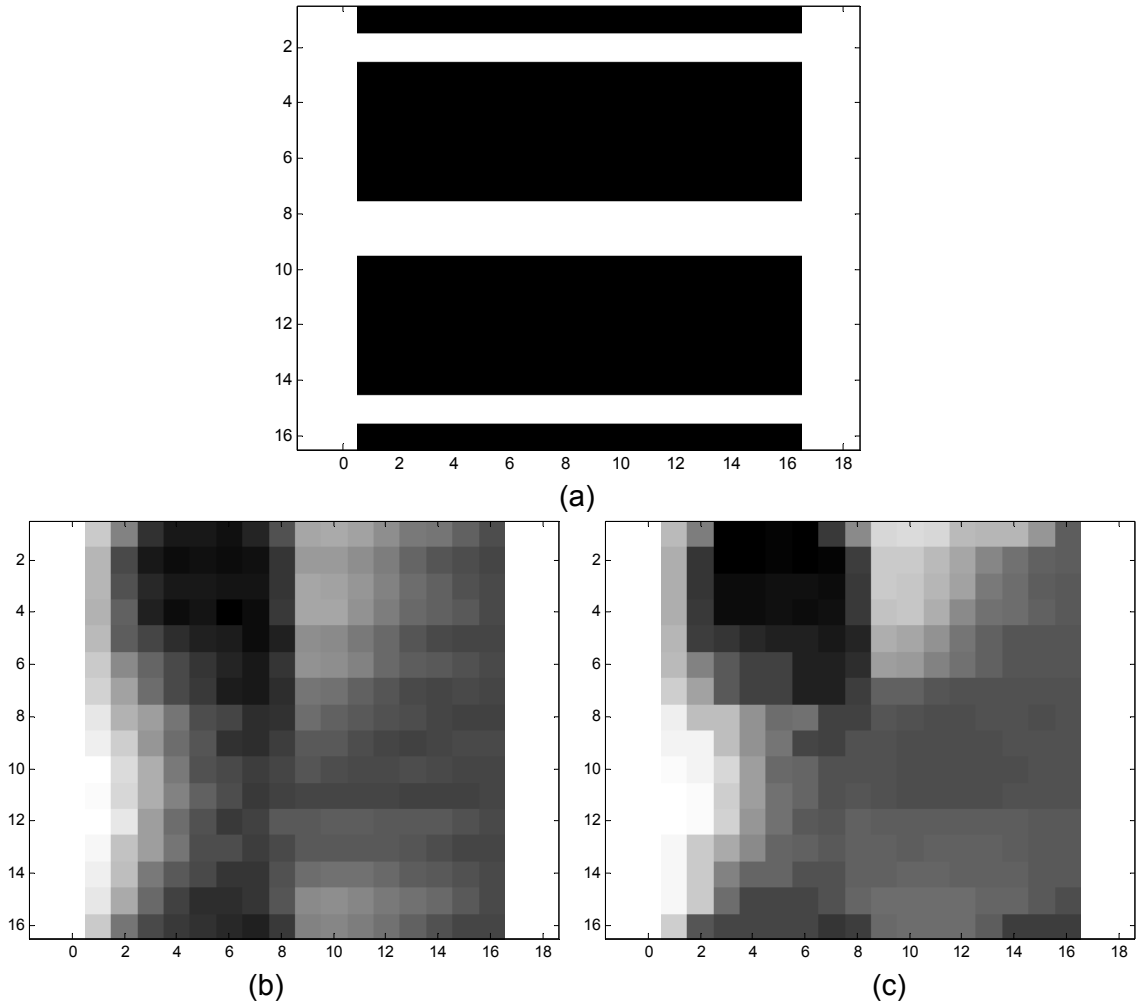

Figure 8: (a) Optimal sampling pattern that includes central phase encode levels, (b) the initial guess generated from a backward projection of the zero-padded Fourier matrix on the undersampled data, and (c) the much smoother reconstructed image.

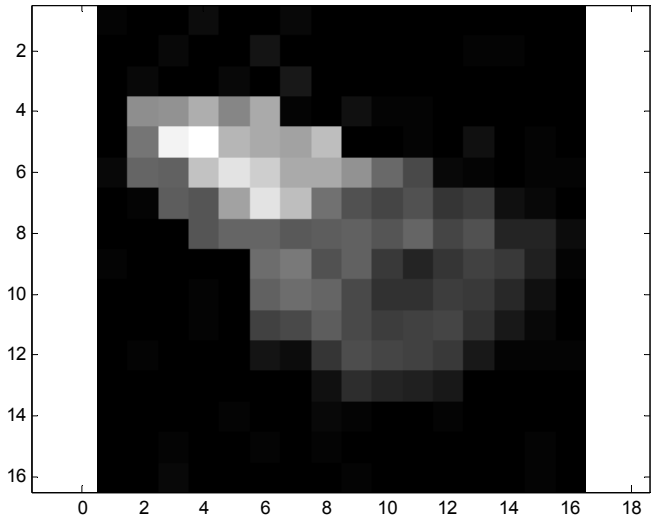

Figure 9: Second test object: Another 16x16 image derived from a web icon.

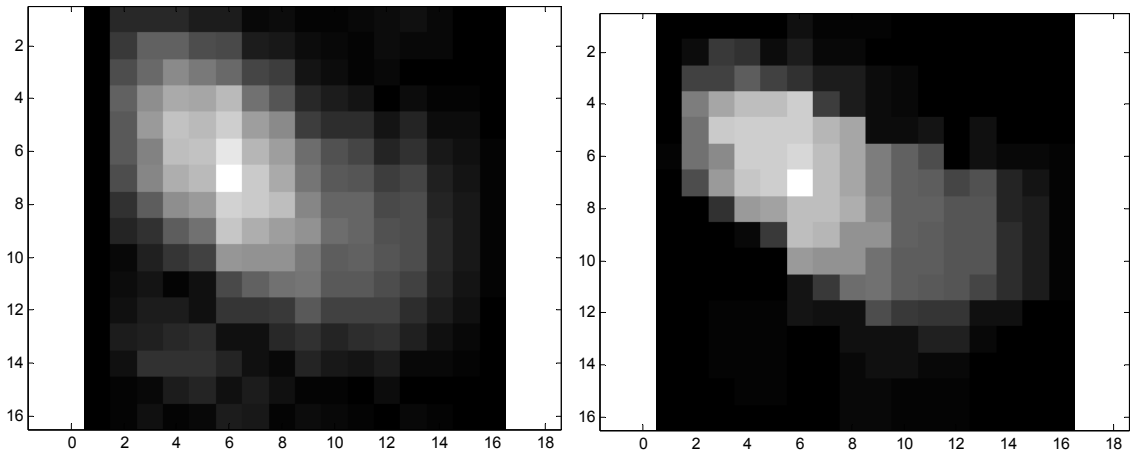

Figure 10: (a) The blurry initial guess generated from a backward projection of the zero-padded Fourier matrix on the undersampled data and (b) the reconstructed image.

Conclusion

The main goals of this project were to develop an understanding of compressed sensing in the context of its application in MRI, to explore optimal sampling, and to develop optimal sampling patterns for small images with 16x16 pixels. The reconstructed images look reasonably good and certainly better than those generated by back-projecting the zero-padded Fourier matrix on the undersampled data. Overall, this project was a great learning experience, and I would like to study ahead and further explore the intricacies of the mathematics behind compressed sensing.

References

[1] R. G. Baraniuk, “Compressive Sensing [Lecture Notes],” IEEE SPM, vol. 24, no. 4, pp. 118-121, July, 2007.
[2] E. Candès and M. B. Wakin, “People Hearing Without Listening: An Introduction To Compressive Sampling,” (To appear in IEEE SPM).
[3] E. Candès, “Compressive sampling,” Proceedings of the International Congress of Mathematicians, Madrid, Spain, 2006.
[4] M. Lustig, D. 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, Dec., 2007.