ECE-435: Project 3 - Wavelets, Compression and Compressed Sensing in Magnetic Resonance Imaging

Zachary Bakhchi

May 15, 2026

1 Introduction

Magnetic Resonance Imaging (MRI) data is highly dense and expensive to store. Transforming images into sparse representations, particularly using wavelets, improves compressibility. Combined with compressed sensing, these methods enable accurate reconstruction from fewer measurements, reducing storage requirements. Analysis was done on data from OpenNeuro.

2 Results

2.1 Navigating The Volumes

Patient CSI1's T1 and T2 anatomical session were considered. In figure 1 is the center index sagittal, coronal, and axial cross sections of the patients brain.

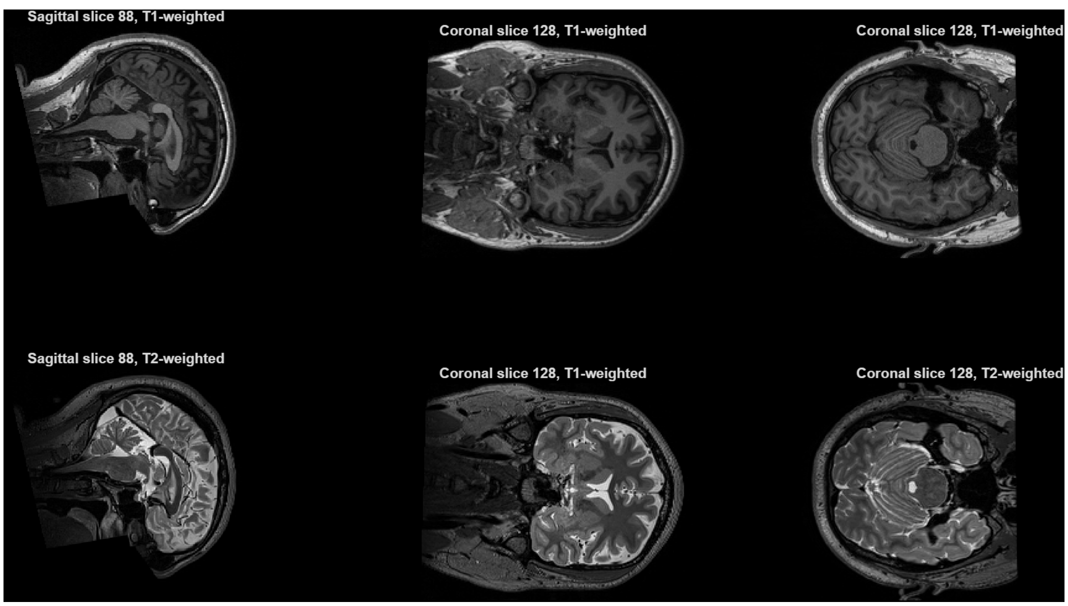

Figure 1: Center Index Sagittal, Coronal and Axial Cross-Sections

Different brain tissue can be made out from the different views. Using the provided table, the bright white of the T1-Weighted images is the white matter while the darker structures are grey matter and cerebrospinal fluids. This is the opposite in the T2-Weighted images case, where the cerebrospinal fluids is now the bright features and the others are now dark. In both sessions grey matter is - not surprisingly - grey. Based on these observations it can be concluded that T1 weighting rewards lower values while T2 rewards higher.

Structure	T1 (ms)	T2 (ms)
White Matter	510	67
Grey Matter	760	77
Cerebrospinal Fluid	2650	280

Table 1: Relaxation Times for Different Brain Tissues

These conclusions can also be supported using the pulse sequence timing parameters which can be obtained with the correspoding .json files: repetition time (TR), echo time (TE), and inversion time (TI). The T1-weighted scan uses a short TE, T2 effects are minimized and a long TI, emphasizing the T1 contrast. On the other hand, the T2-weighted scan uses a longer TE and TR, minimizing the T1 effects. To match the relaxation times given above and realistic MRI acquisition parameters, the units in the given .json files must be in seconds. The parameters, converted to milliseconds, are shown in the table below - matching with standard MRI pulse sequence parameters.

Sequence	TE (ms)	TR (ms)	TI (ms)
T1-Weighted	1.97	2300	900
T2-Weighted	422	3000	N/A

Table 2: MRI Pulse Sequence Parameters

For the subsequent methods applied to the MRI data, we select a guiding image corresponding to the central slice of the sagittal plane of the T1-weighted volume.

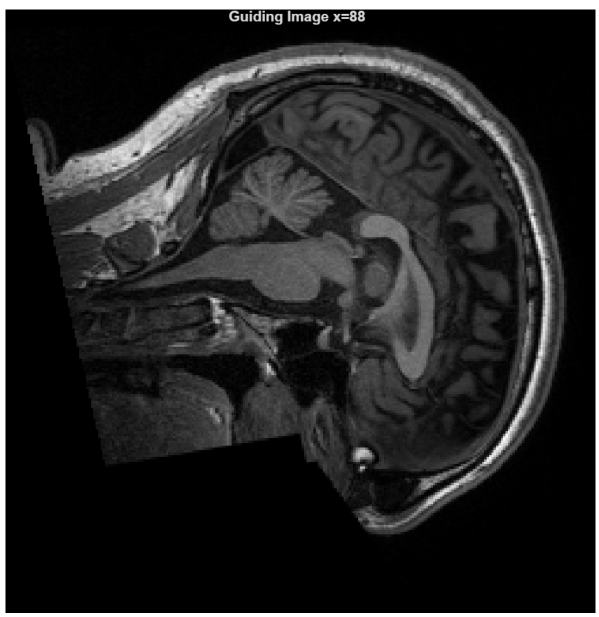

Figure 2: Guiding Image

2.2 Fourier Domain Compression?

As the section title may suggest, what would compression in the Fourier Domain look like? Before that question is answered, the guiding image in the Fourier Domain should be found. This can be seen in Figure 3, but is this image really useful? The sparsity of the image can be measured using a histogram, as seen in Figure 4.

Its pretty obvious that the image is not already sparse in the Fourier Domain, so instead, for multiple percentages s, the lowest-magnitude s percent values of the image can be set to zero to simulate sparsity. For the images below in Figures 5-8, this was done. All, except for Figure 8 where s = 0.99, the images look almost identical. However, looking at the mean squared errors of the image in

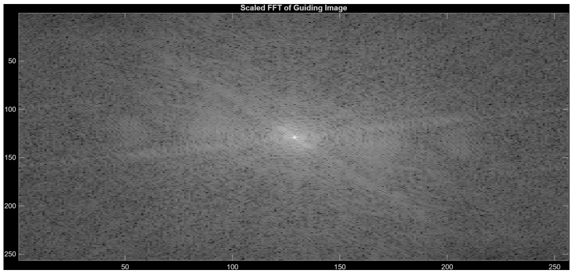

Figure 3: Fourier Domain Image

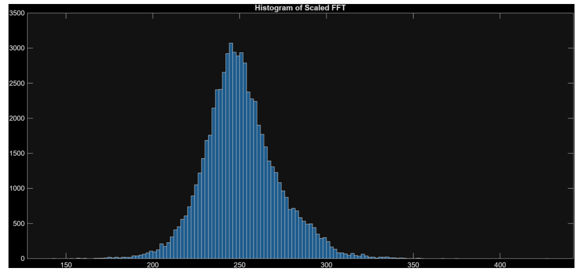

Figure 4: Fourier Domain Image Histogram

Table 2, a different story is told. MSE continues to grow as s increases, and peaks at s = 0.99. While the image may be sparse in the Fourier Domain, this basis of representation will not be good enough for compressed sensing as the phase of the Fourier Domain values, even in low magnitudes, are being lost and is very important to the image and its quality. So, sparsity cannot be achieved in the Fourier Domain, so a new domain should be considered.

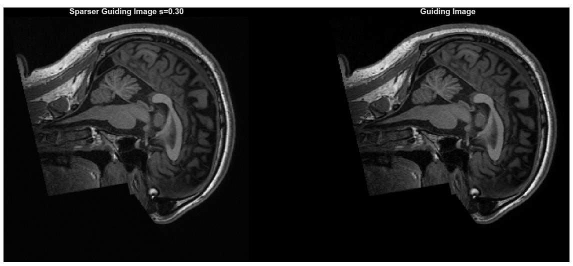

Figure 5: Sparse Fourier Domain Image, s = 0.30

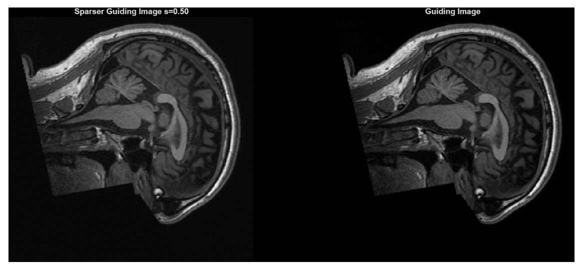

Figure 6: Sparse Fourier Domain Image, s = 0.50

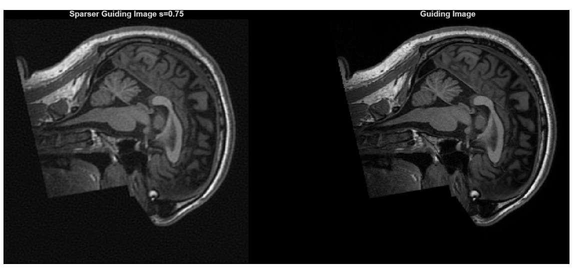

Figure 7: Sparse Fourier Domain Image, s = 0.75

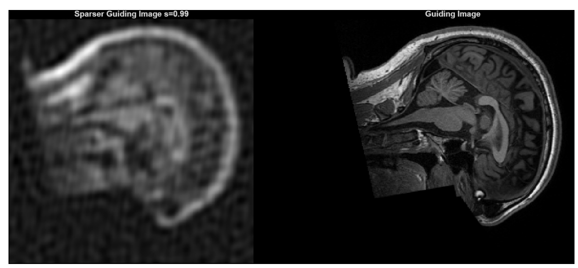

Figure 8: Sparse Fourier Domain Image, s = 0.99

S_Value	**MSE**
0.30	6.7639e-05
0.50	0.00023824
0.75	0.00088236
0.99	0.011052

Table 3: MSE values for different S-values

2.3 Wavelet Domains

The new domain that will be considered will have a wavelet basis. Specifically, the Haar, Daubuchies 4, and Coiflet 3 bases. To familiarize the workings of wavelet transform functions, a 2-level 2D wavelet decomposition was taken on the standard MATLAB 'cameraman' image. The second level approximation is placed in the top left corner, and surrounded by the first level detail components, and then the second level detail components. The detail components correspond to the horizontal, diagonal, and vertical. The transformations on the image are displayed below.

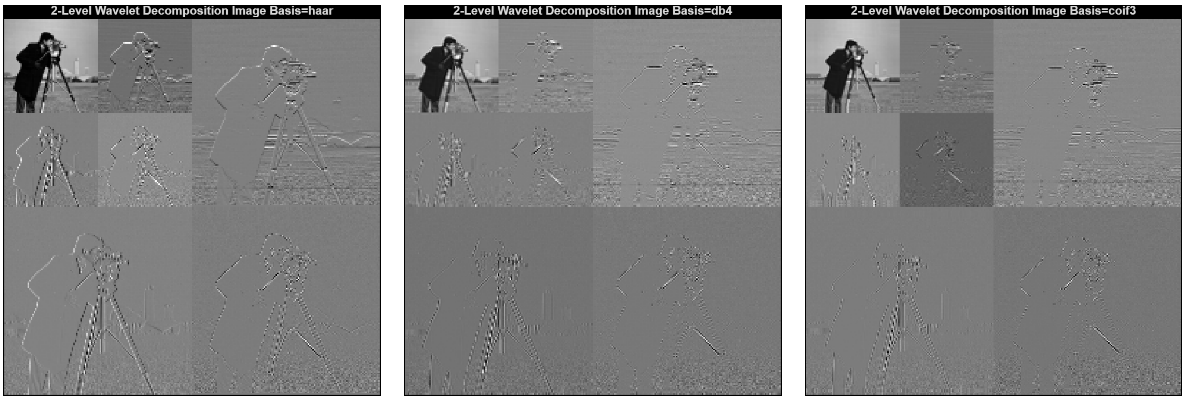

Haar Wavelet Basis Daubuchies 4 Wavelet Basis Coiflet 3 Wavelet Basis

Figure 9: Comparison of Wavelet Bases

Immediately, differences between the transformations can be observed. Haar is more noticeably blocky compared to the others and Daubuchies 4 and Coiflet 3 are more similar to each other. Lets test these bases on a standard basis element - a 256 x 256 matrix where all elements are zero and the middle element is one.

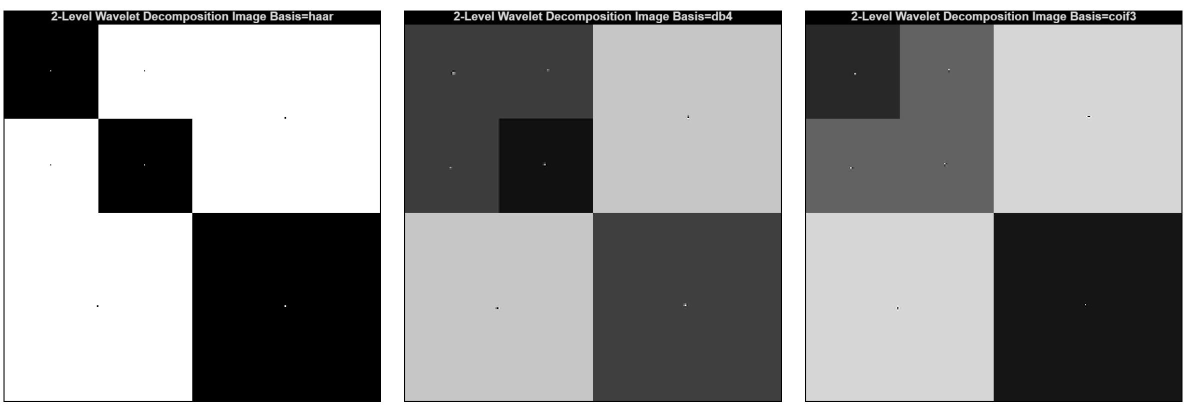

Haar Wavelet Basis Daubuchies 4 Wavelet Basis Coiflet 3 Wavelet Basis

Figure 10: Comparison of Wavelet Bases on Standard Element

Like the cameraman image, similar conclusions can be drawn about the bases of the wavelet transforms applied to the standard element. Haar is blocky and localized while Daubuchies 4 and Coiflet 3 are more spread out. However, now the two can be differentiated better. The two spread out bases more smoothly, but differently as seen by the locations of dark and light grays throughout the images. This can be further seen by taking the magnitude of the Fourier Transforms of these images in the figure below.

As expected, Haar is blocky and now its seen that across spatial frequency that Daubuchies 4 and Coiflet 3 wavelet bases are spread out differently. Daubuchies 4 emphasizes energy in the middle, while Coiflet 3 has peaks that start high and then decay to zero. These observations match the canonical representations of the wavelets. Haar exhibits sudden discontinuities, while Daubechies and Coiflet

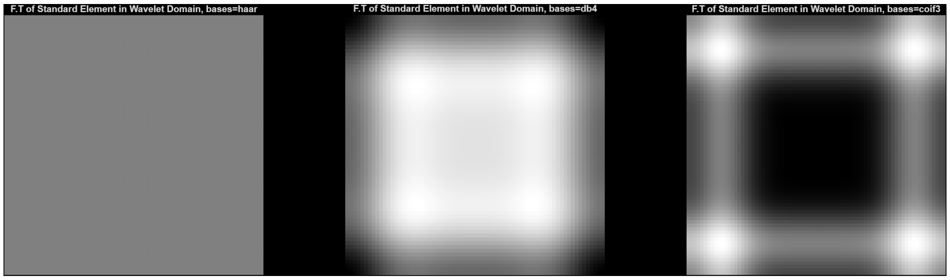

Figure 11: F.T. of Wavelet Domain Standard Element

wavelets are more continuous and have different energy localities. Now that more is known about these wavelets, they can be applied to the guiding image. The 2-level decompositions of the guiding image in these three bases can be seen in the figures below.

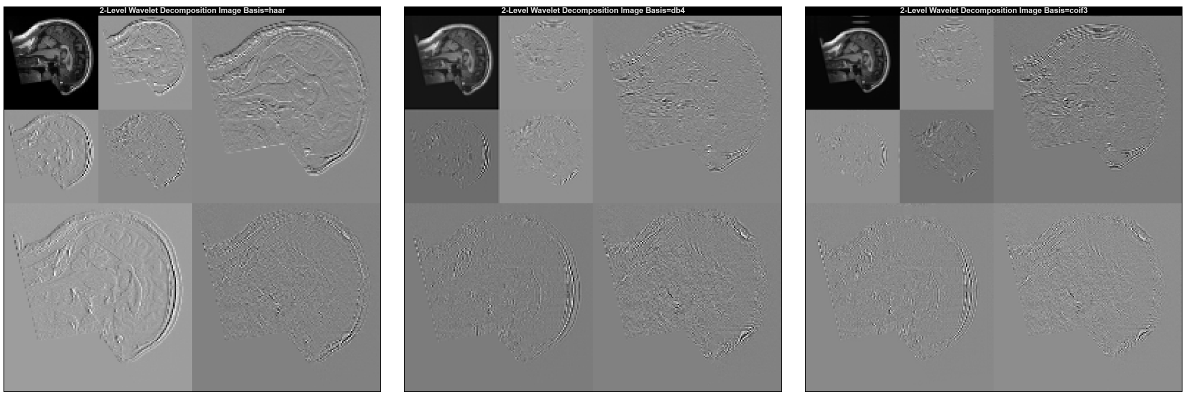

Haar Wavelet Basis | Daubuchies 4 Wavelet Basis | Coiflet 3 Wavelet Basis

Figure 12: Comparison of Wavelet Bases on Guiding Image

Based on what's been seen already with the cameraman, the standard element and its Fourier Transform, these wavelet decompositions make sense. However, the original question still remains; is there sparsity in this domain? By looking at the histograms of the 1-level, 2-level, and 3-level decompositions of these bases, the question can be answered. These histograms can be seen in the figures below.

Immediately, it is observed that the 2-level and 3-level bases are more sparse than the rest. And based on what was learned already about these bases, Haar should not be used. The 3-level Daubechies 4 wavelet basis is probably the best candidate for compressed sensing. This can be further tested by zeroing out the lowest 10 percent of wavelet coefficients, reconstructing it, and calculating its MSE. This was done, in the figure below and it can be concluded that this wavelet domain is sparse with an MSE of approximately 0.

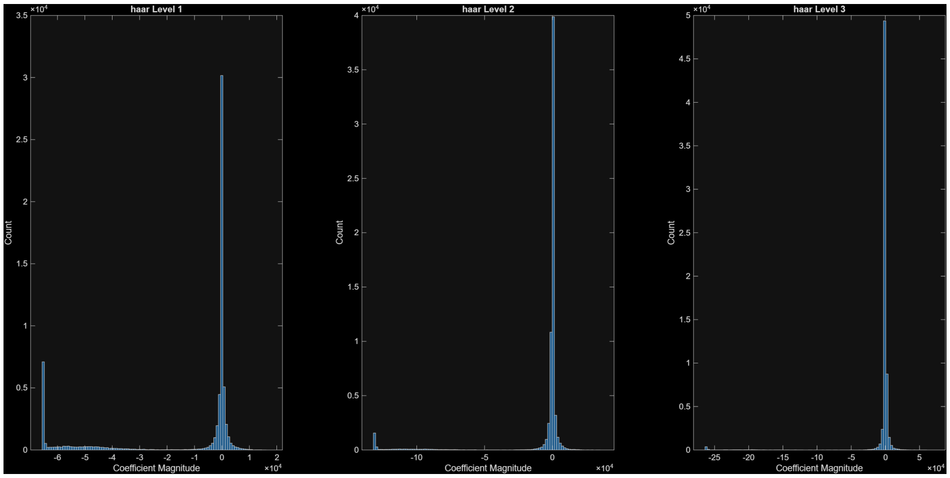

Figure 13: Haar Wavelet Basis Histogram

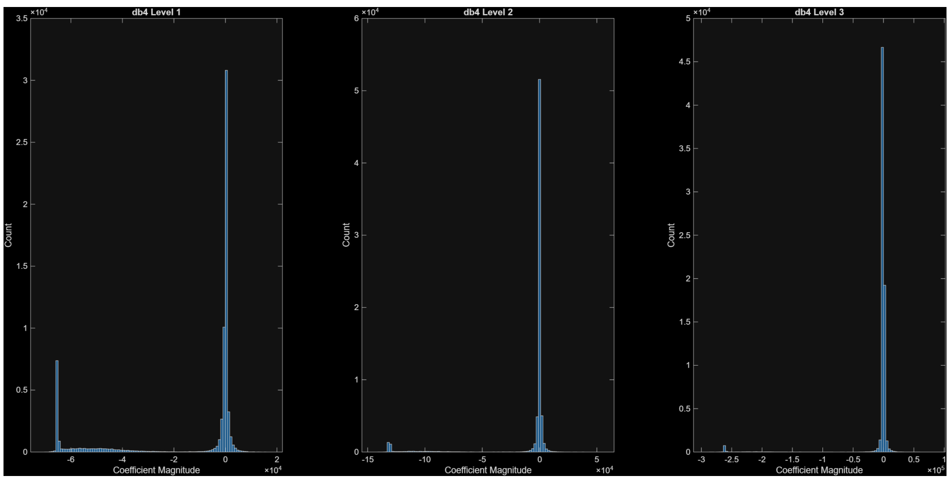

Figure 14: Daubechies 4 Wavelet Basis Histogram

2.4 How Sparse?

As the title of this section suggests, the new question being asked is how sparse can these domains get. To test, the MSE between the reconstructed image and the guiding image can be plotted as a function of s, the percent of coefficients being deleted. This was done on the 2-level and 3-level bases that have been mentioned already, with the result in the figure below.

These results are phenomenal - for each bases error remains approximately zero until s reaches 90 percent. In fact, for the 3-level Daubechies 4 wavelet basis that will be used for compressed sensing, up to 98 percent of coefficients can be deleted before MSE reaches 10 percent. However, it is important to note that once that threshold is reached, then the error is much more obvious - especially since the MSE quickly approaches 103 percent when all coefficients are deleted. This can be further demonstrated by looking at the reconstructed images for multiple s values for the different bases on the guiding image and other images (Figures 18-23).

Each image, guiding and T2-weighted in the Daubechies 4 case, is able to maintain the same look for each s value, that is until s = 0.95, where differences are seen. As expected, each 2-Level wavelet reconstruction struggles to be reconstructed because of their lack of sparsity. However, the 3-Level wavelet reconstructions are an improvement from the 2-Level, as the brain could still be made out even though there is some error.

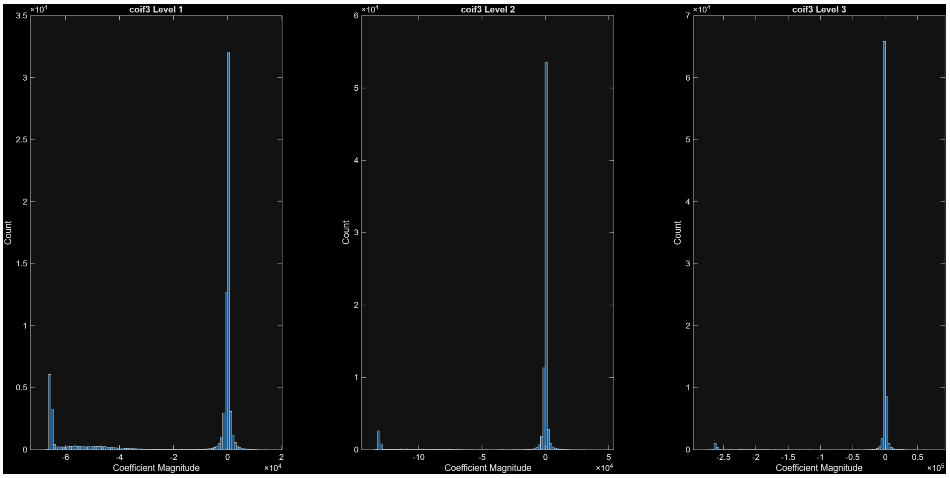

Figure 15: Coiflet 3 Wavelet Basis Histogram

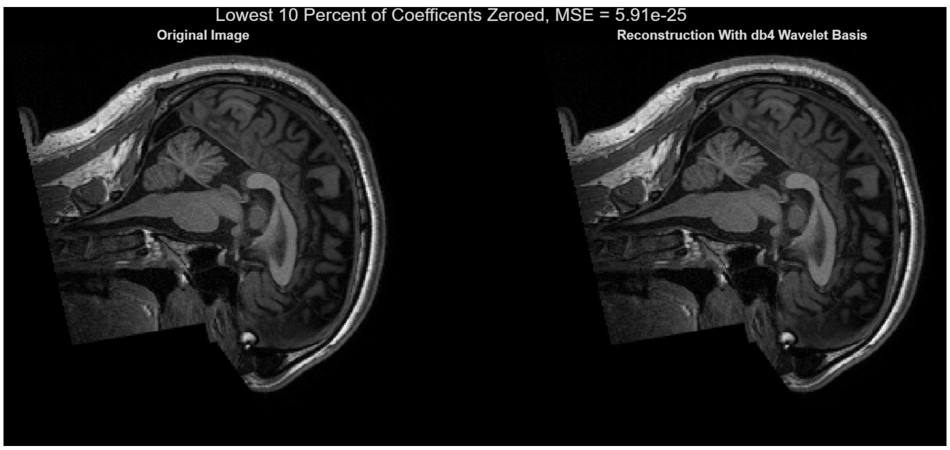

Figure 16: Daubechies 4 Wavelet Reconstruction

2.5 Compressed Sensing

Compressed Sensing is what the previous sections have built upon, helping reconstruct spatial frequency domain MRI images into the best representation of the ground truth. While MRI images are not sparse in its pixel (image) domain, they become sparse after being transformed into an appropriate basis, in this case the 3-level Daubechies 4 wavelet basis - which was very sparse as seen before.

To simulate realistic MRI acquisition while utilizing compressed sensing techniques, the Fourier transform of the guiding image was taken as the fully sampled ground truth k-space data. This k-space representation was then undersampled using a Bernoulli random sensing mask, where measurements were randomly retained with probability p and discarded otherwise. This process simulates accelerated MRI acquisition by reducing the number of collected Fourier measurements while still attempting to reconstruct the original image from incomplete data. Now its just an optimization problem using the Iterative Shrinkage-Thresholding Algorithm (ISTA). For varying values of the sparsity parameter λ and iterations of ISTA, compressed sensing was performed for different values of p. In addition, convergence behavior was analyzed over 50 independent trials, where a different randomly generated undersampling matrix was used for each trial. The average convergence was plotted over each loss function.

The results from compressed sensing can be seen below. (Figures 24-27)

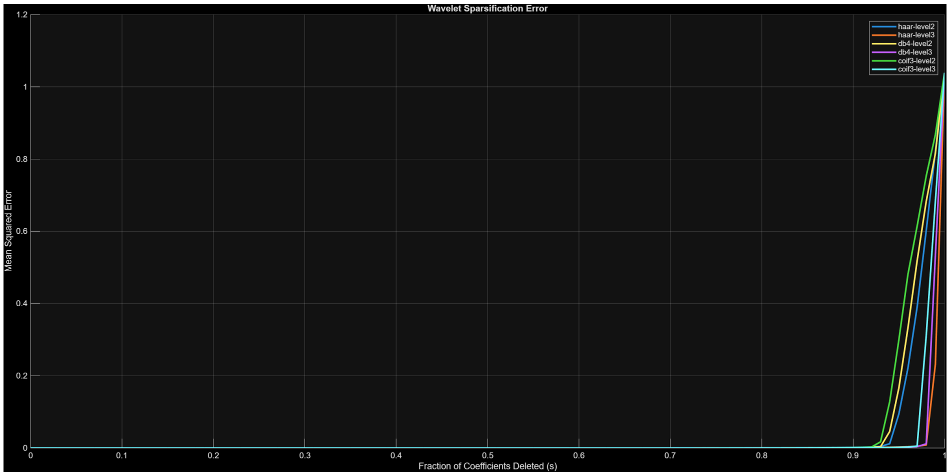

Figure 17: MSE vs. s

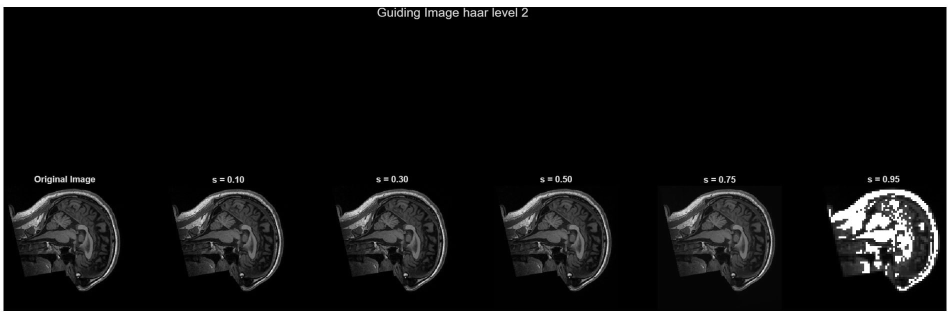

Figure 18: Reconstructed Image, 2-Level Haar

The p values chosen were [0.3, 0.5, 0.7, 0.9]. Each p value used a different value λ and different number of iterations to get the best reconstructed image, except for p = 0.9, where the parameters did not matter because the image was too undersampled and could not be reconstructed with enough detail. On the other hand, at p = 0.3 the image was able to be reconstructed with $\lambda = 30$ and only 500 iterations. At p = 0.5, the image beocmes blurry, but mostly able to maintain most detail. In this case, it was found helpful to increase λ to 40 instead of 30. Finally, at p = 0.7, the image is also able to be reconstructed but with the most blur due to the high about of error from the undersampled image. In this case it was useful to decrease λ to 25 and increase the number of iterations to 1000. In each case of compressed sensing, the loss function was able to converge to a valid, and best case, reconstruction.

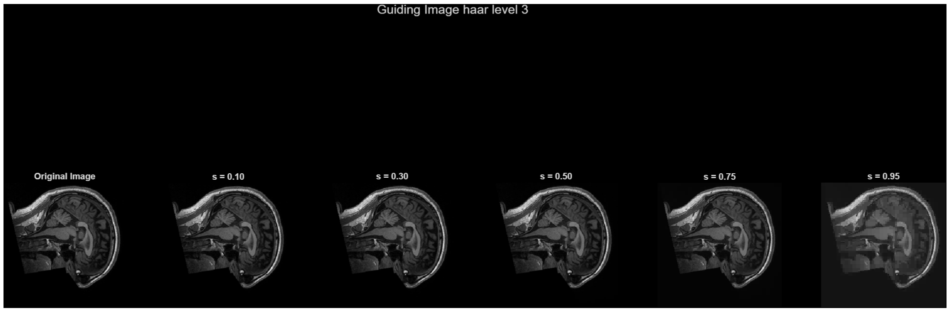

Figure 19: Reconstructed Image, 3-Level Haar

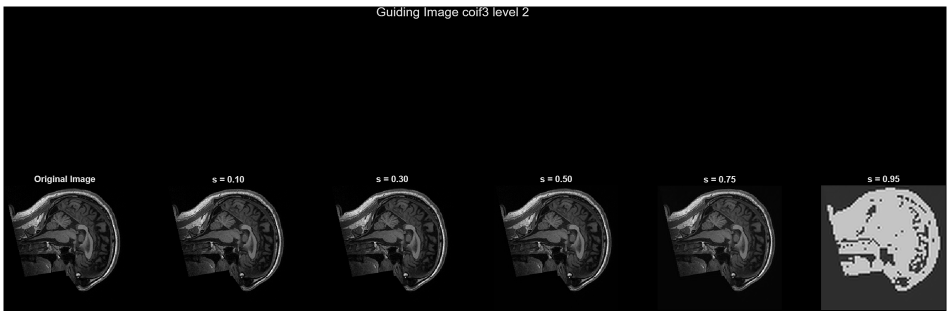

Figure 20: Reconstructed Image, 2-Level Coiflet 3

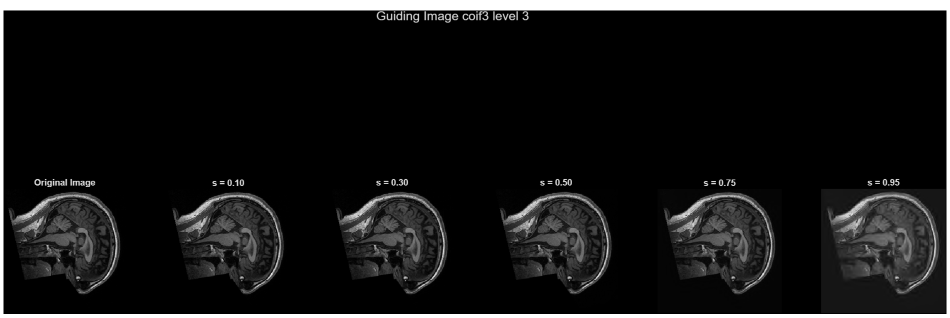

Figure 21: Reconstructed Image, 3-Level Coiflet 3

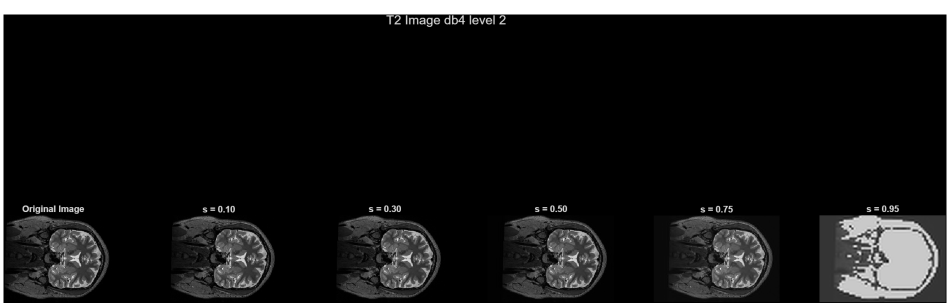

Figure 22: Reconstructed Image, 2-Level Daubechies 4

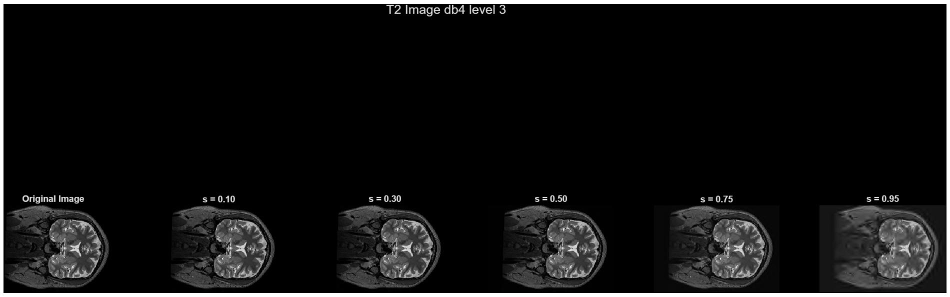

Figure 23: Reconstructed Image, 3-Level Daubechies 4

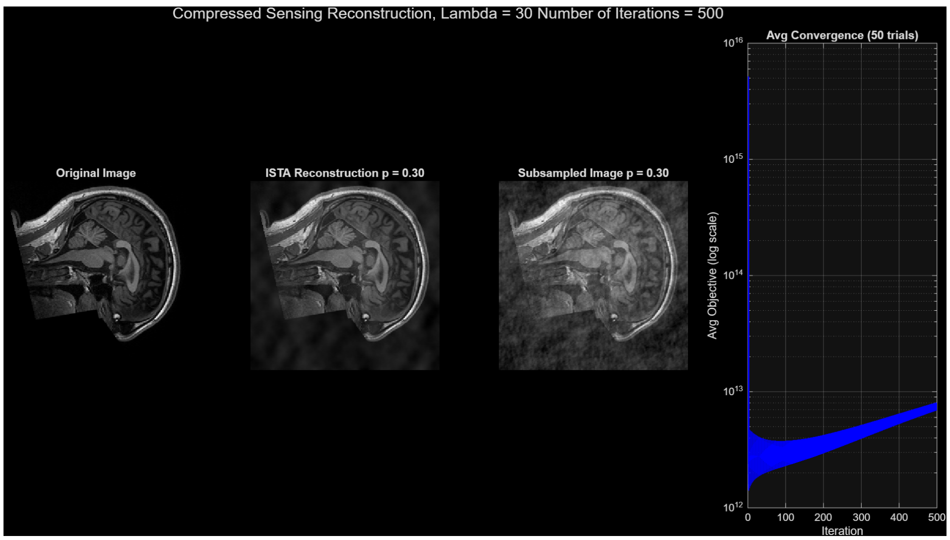

Figure 24: Compressed Sensing, p = 0.30

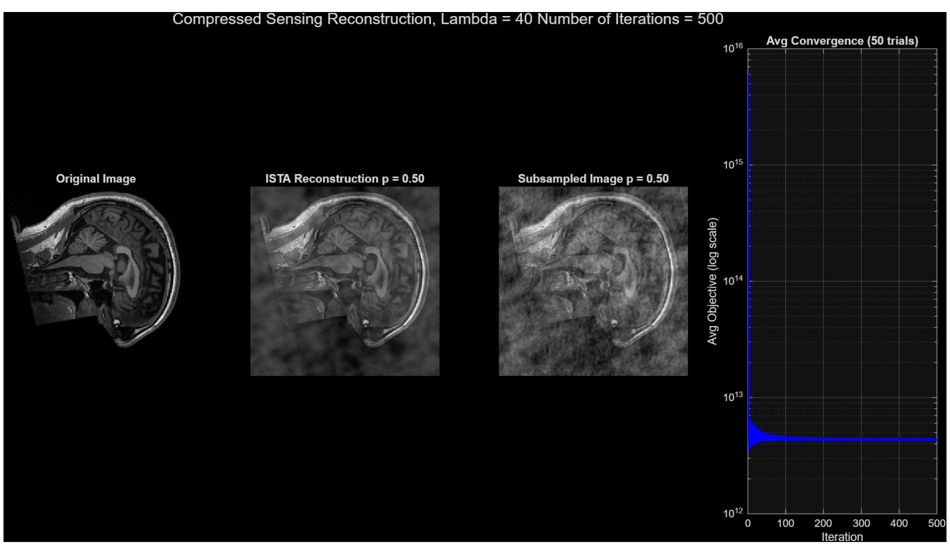

Figure 25: Compressed Sensing, p = 0.50

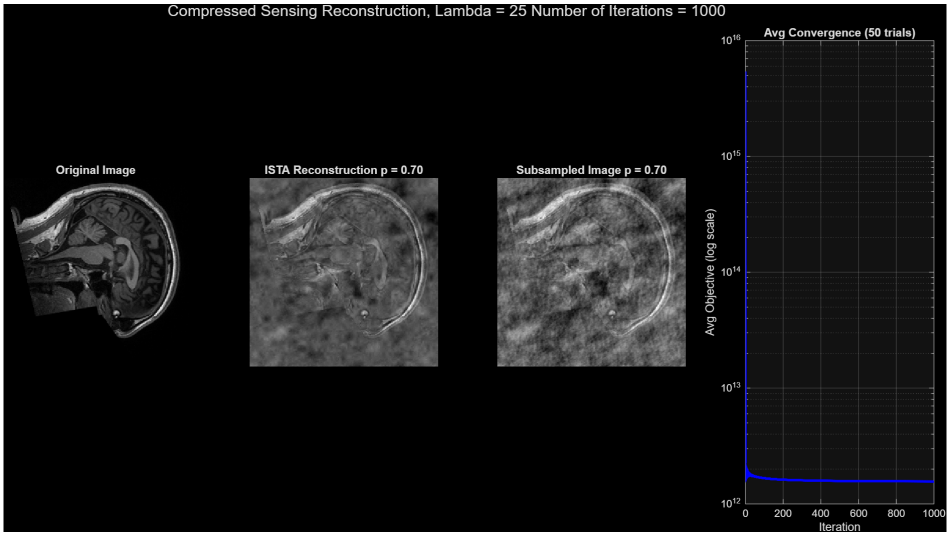

Figure 26: Compressed Sensing, p = 0.70

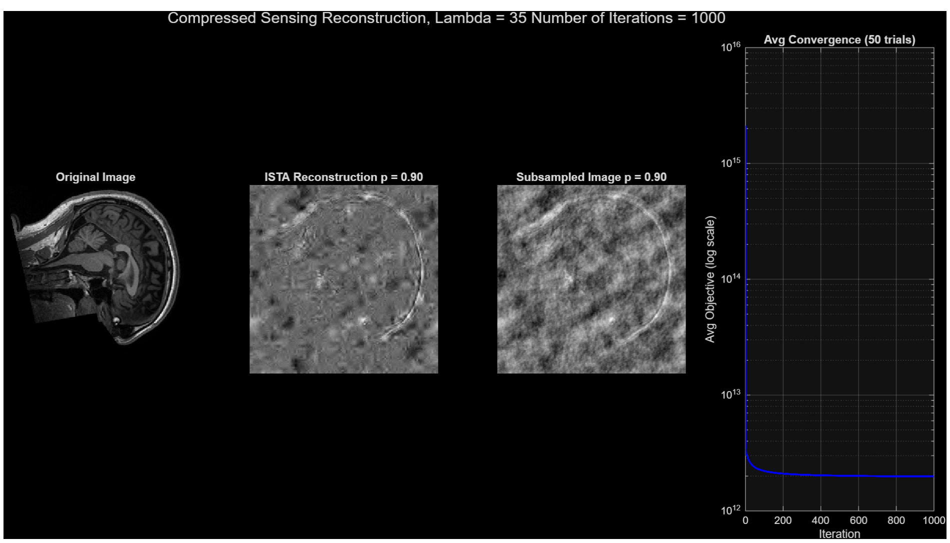

Figure 27: Compressed Sensing, p = 0.90