Subsampling and reconstruction of bandlimited images with universal sampling sets

Leighton P. Barnes
Department of Electrical Engineering
Stanford University
Email: lpb@stanford.edu

***Abstract*—We investigate the subsampling and reconstruction of bandlimited images at universal sampling sets. Theoretically such sampling sets should guarantee the reconstruction of an image that is k-sparse in the DFT domain with only k samples. We find that, due to matrix conditioning issues, more than k samples are generally required, and we compare the reconstruction results to those from the sparse recovery algorithms IHT and CoSaMP. This method achieves lower MSE for a given sampling rate but it requires knowledge of the support of the sparse image.**

I. Introduction

Reconstructing signals from a limited set of subsamples is a fundamental problem in signal processing. Sometimes it is done for simple compression or computation purposes. It is often easier to work with smaller, reduced data sets. Other times, particularly in sensing applications, only a limited number of samples is available. In this case we need to exploit the structure of the signals in order to fill in the missing data.

In the past decade there has been an explosion of work related to subsampling by using the methods from compressed sensing. These methods assume that the signals under consideration are sparse in some known basis and then use iterative algorithms to find the support of the signal (in that sparse basis) and the associated coefficients. The 'subsamples' for these methods are often taken with a random sensing matrix (i.e. with i.i.d. Gaussian entries or random harmonic frames). There are theoretical guarantees with these sensing matrices that reconstruction will fail only with exponentially low probability.

In contrast with the iterative methods from compressed sensing, a purely linear approach can be taken to signal reconstruction if the support of the sparse signal is known. In this genie-provided-support case, the reconstruction of the signal relies only on the invertability of the submatrices of the sensing matrix. Recent mathematical results on the invertability of submatrices of the DFT matrix [1],[3] suggest that if the samples of a signal that is sparse in the DFT domain are taken at specific indices which correspond to a universal sampling set (USS), then perfect reconstruction should always be possible. This is a deterministic guarantee as opposed to the probabilistic ones from random matrices.

Motivated by these theoretical results, this project was an experiment to see if sampling at the locations of a universal sampling set results in a good reconstruction of a bandlimited image (i.e., an image which is sparse in the DFT domain). To this end, we attempt the subsampling and subsequent reconstruction of example images with a universal sampling set. We then compare these results against the popular sparse reconstruction algorithms of iterative hard thresholing (IHT) and compressive sensing matching pursuit (CoSaMP). Our comparison focuses on what information is required to perform the reconstructions, for what rates the reconstructions can succeed, and the visual quality of the reconstructions.

II. Background

A. Universal Sampling Sets

Let's first consider one-dimensional, discrete-time signals $x[n]$ of length N. We will later show how this case can be related to the two-dimensional signals which arise in the image processing context. In order to construct universal sampling sets for these signals, we will be concerned with the prime factors of N. Let $N = N_1 N_2 \dots N_p$ be its prime factorization.

Suppose we have access to the samples of $x[n]$ taken at indices $n = m_i$ where $m_i \in \{1, 2, \dots, N\}$ and $i = 1, \dots, M$. Call $I = \{m_i\}$ the set of M distinct sampling indices. Suppose further that $x[n]$ is k-sparse in the DFT domain with support corresponding to the indices J. In other words,

$$x[n] = \frac{1}{N} \sum_{i \in J} a_i e^{j \frac{2\pi}{N} in} \ . \tag{1}$$

We define the matrix E_I to be the M by N submatrix of the identity which selects out precisely the samples in I. This gives

$$E_I x = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 & 0 \\ \vdots & & & & & \vdots \\ 0 & \dots & 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} x[1] \\ \vdots \\ x[N] \end{bmatrix} = \begin{bmatrix} x[m_1] \\ \vdots \\ x[m_M] \end{bmatrix} \ . \tag{2}$$

Let F be the N by N DFT matrix. Rewriting (1) as a matrix relationship yields

$$\begin{aligned} x &= F^{-1} E_J^T a & (3) \\ \implies E_I x &= E_I F^{-1} E_J^T a & (4) \\ \implies a &= (E_I F^{-1} E_J^T)^{-1} E_I x & (5) \\ \implies x &= F^{-1} E_J^T (E_I F^{-1} E_J^T)^{-1} E_I x \ . & (6) \end{aligned}$$

Thus if $E_I F^{-1} E_J^T$ is invertible, we can reconstruct all of x from the subsamples $E_I x$. Note that $E_I F^{-1} E_J^T$ is simply the submatrix of the inverse DFT matrix with rows I and columns J. This leads to the following definition. A sampling set I is a *universal sampling set* if $E_I F^{-1} E_J^T$ is invertible for any choice J of M different frequency components.

There is a complete characterizations of what the universal sampling sets look like if N is a prime power [1]. It is an open problem to find such a characterization for general N, but there are known examples which are easy to construct. For example, Proposition 7 from [1] shows that if the m_i are integers that are evenly spaced distance Δ apart (and then taken modulo N), then I will form a universal sampling set provided that Δ is coprime to the N_i.

B. Vectorization

We now turn to the issue of how this framework can be applied to two-dimensional signals such as images by vectorization with a specific order. We need to ensure that if an image has a k-sparse two-dimensional DFT, then once it is vectorized we can still formulate the reconstruction as in (6) above. Let $x[n,m]$ be an $\mathcal{N}_1$ by $\mathcal{N}_2$ image with $\mathcal{N}_1, \mathcal{N}_2$ coprime and let $N = \mathcal{N}_1 \mathcal{N}_2$.

If $x[n,m]$ is vectorized in the usual lexiconographic order, then the corresponding two-dimensional DFT matrix for operating on the vectorized image becomes the Kronecker product $F_{\mathcal{N}} = F_{\mathcal{N}_1} \otimes F_{\mathcal{N}_2}$ [4]. Lemma 2 from [4] shows that the usual N by N DFT matrix F_N is the same as $F_{\mathcal{N}}$ up to row and column permutations. The Lemma further shows that $E_I F_{\mathcal{N}}^{-1} E_J^T = E_{c(I)} F_N^{-1} E_{cb^{-1}(J)}^T$ where

$$c^{-1}(k) = (k \bmod \mathcal{N}_1, k \bmod \mathcal{N}_2)$$

and b is a specific bijection. Because $\mathcal{N}_1$ and $\mathcal{N}_2$ are coprime, the Chinese remainder theorem ensures that c is well-defined and is also a bijection.

The key thing here is that if we construct a universal sampling set I for a one-dimensional signal of size N, then $c^{-1}(I)$ gives a corresponding set of sampling indices for the image.

C. IHT/CoSaMP

Both IHT and CoSaMP are iterative algorithms for finding k-sparse reconstructions of a signal x given the measurement $y = \Phi x$ where Φ is an M by N measurement system and $M \leq N$. The problem is to find $a = B^{-1}x$ for an orthonormal basis B such that $\|a\|_0 \leq k$ (i.e. a is k-sparse).

IHT works essentially by gradient descent. It makes an estimate $\hat{a}$ which is updated iteratively by

$$\hat{a} = T_k(\hat{a} + F^H(y - F\hat{a}))$$

where $F = \Phi B$ and T_k truncates to the k coefficients with highest magnitude. The $F^H(y - F\hat{a}))$ term acts as the gradient and T_k enforces that the estimates are k-sparse. CoSaMP is similar but it works by keeping an estimate $\hat{I}$ of the support of a and then exploring which other components might be needed based on the residue $y - F\hat{a}$. For a full description of the CoSaMP algorithm see [2].

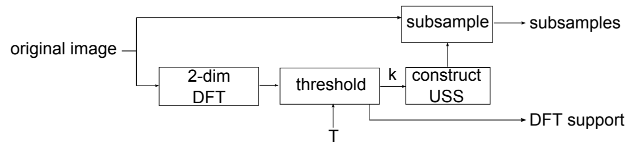

Fig. 1: Block diagram of the subsampling procedure.

III. Methods

In order to test this universal sampling set framework with real images I used the following procedure. To subsample the images I followed the process depicted in the block diagram in Figure 1. First, the two-dimensional DFT of the original image was found. The DFT coefficients were thresholded to decide what k should be and which k coefficients would provide the best sparse approximation of the image. The locations of those coefficients were recorded so they could later be used for the reconstruction. Once k was known, a universal sampling set was constructed according to k and the dimensions of the image. I used the sampling sets $c^{-1}(I)$ for an $\mathcal{N}_1$ by $\mathcal{N}_2$ image where $I = \{i\Delta \bmod N\}$ and $i = 1, \ldots, k$. Finally, the subsamples of the image at the indices $c^{-1}(I)$, as well as the support of the DFT coefficients, were provided to the reconstruction algorithm.

The reconstruction algorithm amounts to setting up the system of equations in (4) and solving it. To do this, vectorized versions of x and F need to be constructed. If $M = k$, the number of samples is equal to the sparsity of the image and the system can be solved with a usual matrix inverse. If $M > k$, then we have more samples than should be strictly necessary and the system is overdetermined. This can be done to combat poor conditioning in the submatrices and a pseudoinverse is used to solve the system.

Subsampling for the IHT and CoSaMP trials was done with i.i.d Gaussian sensing matrices Φ with entires $\sim N(0, \frac{1}{M})$. The basis B was the vectorized inverse DFT matrix (scaled to be unitary) so that we are considering the signal to be sparse in the same basis across the different trials.

IV. Experimental Results

I tested these methods on the two images in Figure 2 for various sampling rates M/N both with and without added noise. After resizing the images they were of dimension 85 by 127 and 149 by 81, respectively. Both of these pairs are coprime. Throughout I used $\Delta = 32$ to construct the universal sampling sets which is coprime with all of the dimensions. Example USS reconstructions for $k = M$ and $k = 5M$ can be seen in Figures 3 and 4, respectively. The sampling pattern corresponding to the $k = 5M$ universal sampling set can be seen in Figure 5. For comparison, example CoSaMP reconstructions when $k = 5M$ can be seen in Figure 6.

In Figure 3 we see that when $k = M$, even though there is a theoretical guarantee that the submatrix will be invertible, the reconstruction fails due to poor conditioning of the submatrix. Because of this we are forced to go to higher sampling rates and it becomes more of a question of the properties of

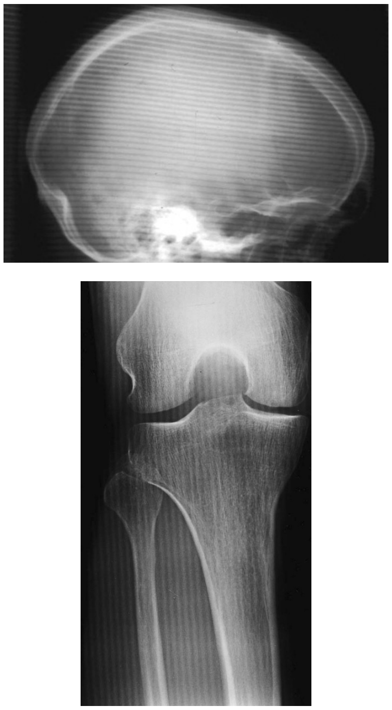

Fig. 2: First (top) and second (bottom) test images thanks to EE368 HW4!

deterministic harmonic frames (in fact, this is how [3] presents this theory).

A plot of the MSE between the original and reconstructed images for the USS method as well as IHT/CoSaMP can be seen in Figures 7 and 8. In Figure 7 we consider the noiseless case and in Figure 8 some variance 10 Gaussian noise is added. Over the range of sampling rates that were tested, IHT either doesn't converge or is just starting to converge. CoSaMP, which is known to work over a wider range of sampling rates, converges in both the noise and no-noise cases for sampling rates greater than approximately 0.1. The USS method gives a good reconstruction for similar rates, and achieves a lower MSE for higher rates. This lower MSE corresponds to improved visual quality as can be seen by comparing Figures 4 and 6.

V. Conclusions

We have demonstrated the use of universal sampling sets in subsampling and reconstructing images which are relatively sparse in the DFT domain. This method achieves lower MSE than sampling with random Gaussian matrices and iterative reconstruction via CoSaMP or IHT, but it requires knowledge of the support of the sparse coefficients. Unfortunately, even though there is a theoretical guarantee that a k-sparse image should be recoverable with k samples taken at the locations given by a universal sample set, poor conditioning of the matrices can prevent such reconstruction.

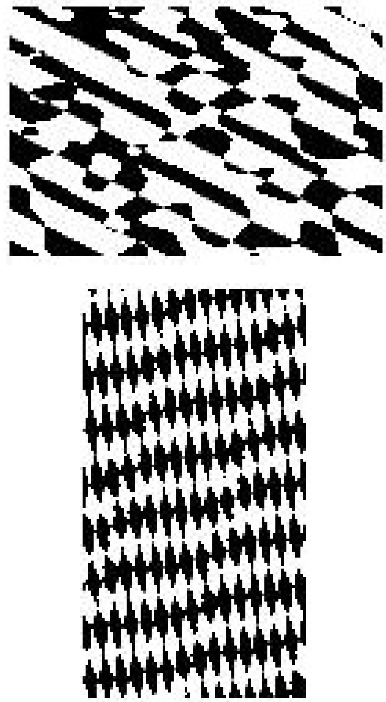

Fig. 3: Failed USS reconstruction when $k = M$. This was using a threshold $T = 5000$ which led to k values of 337 and 355, respectively. In both images this corresponds to $M/N \approx 0.03$.

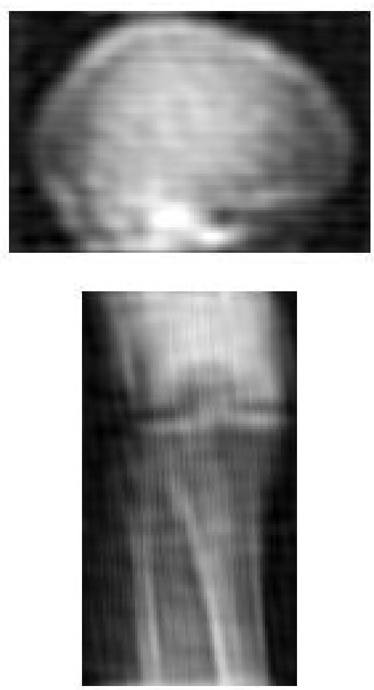

Fig. 4: USS reconstruction when $M = 5k$. This was using a threshold $T = 5000$ which led to k values of 337 and 355, respectively. In both images this corresponds to $M/N \approx 0.15$.

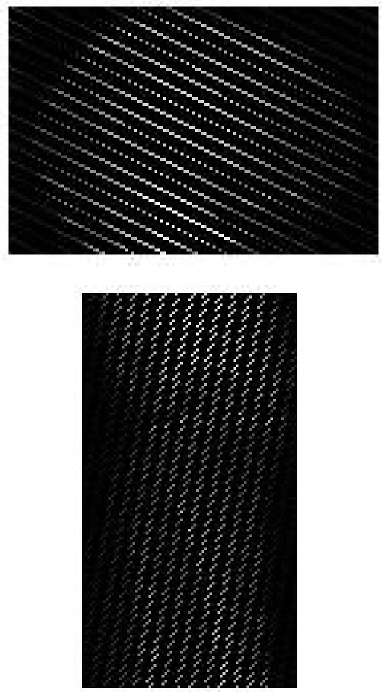

Fig. 5: USS sampling pattern with $\Delta = 32$ when $k = 5M$ for a sampling rate of $M/N \approx 0.15$.

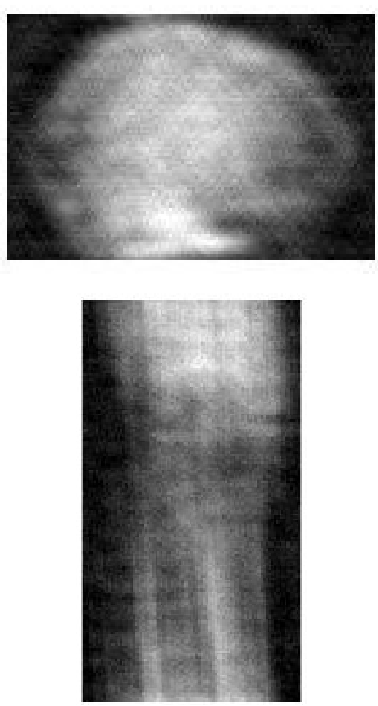

Fig. 6: CoSaMP reconstruction when $M = 5k$ where k was taken as in Figure 4. In both images this corresponds to $M/N \approx 0.15$.

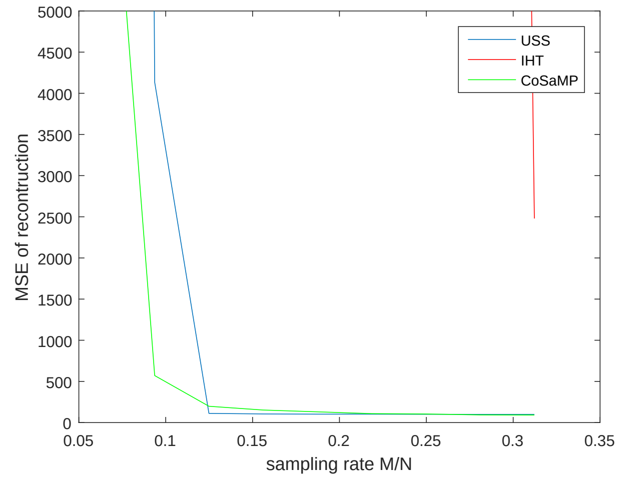

Fig. 7: MSE as a function of sampling rate with no added noise.

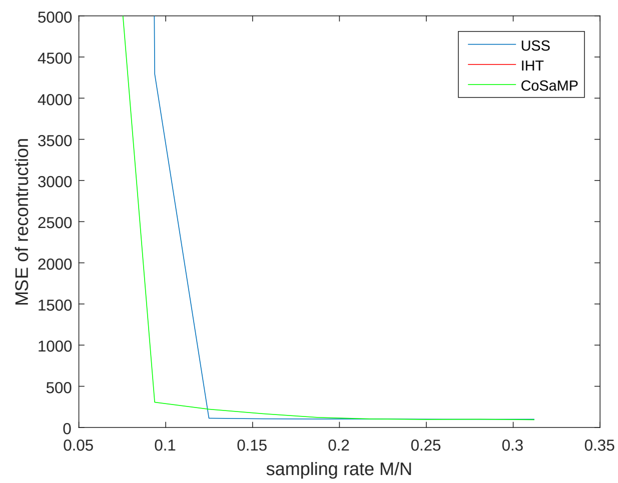

Fig. 8: MSE as a function of sampling rate with added i.i.d. noise added $\sim N(0, 10)$.

Acknowledgments

I would like to thank Kushagr Gupta and Gordon Wetzstein for their work in EE368 and for their helpful comments during the poster session.

References

[1] B. Osgood, A. Siripuram, and W. Wu, *Discrete Sampling and Interpolation: Universal Sampling Sets for Discrete Bandlimited Spaces*, IEEE Transactions on Information Theory 58 (2012), 4176-4200.

[2] D. Needell, and J. A. Topp, *CoSaMP: Iterative Signal Recovery from Incomplete and Inaccurate Samples*, (Accessed 12/3/2015) available at http://arxiv.org/pdf/0803.2392v2.pdf .

[3] B. Alexeev, J. Cahill, and D. G. Mixon, *Full Spark Frames*, The Journal of Fourier Analysis and Applications 18 (2012), 1167-1194.

[4] B. Osgood, *Higher Dimensional Discrete Fourier Transforms*, unpublished.