

Progressive versus Random Projections for Compressive Capture of Images, Lightfields and Higher Dimensional Visual Signals

Rohit Pandharkar
MIT Media Lab
75 Amherst St, Cambridge, MA

Ashok Veeraraghavan
MERL
201 Broadway, Cambridge MA

Ramesh Raskar
MIT Media Lab
75 Amherst St, Cambridge, MA

Abstract

Computational photography involves sophisticated capture methods. A new trend is to capture projection of higher dimensional visual signals such as videos, multi-spectral data and lightfields on lower dimensional sensors. Carefully designed capture methods exploit the sparsity of the underlying signal in a transformed domain to reduce the number of measurements and use an appropriate reconstruction method. Traditional progressive methods may capture successively more detail using a sequence of simple projection basis, such as DCT or wavelets and employ straightforward backprojection for reconstruction. Randomized projection methods do not use any specific sequence and use L_0 minimization for reconstruction. In this paper, we analyze the statistical properties of natural images, videos, multi-spectral data and light-fields and compare the effectiveness of progressive and random projections. We define effectiveness by plotting reconstruction SNR against compression factor. The key idea is a procedure to measure best-case effectiveness that is fast, independent of specific hardware and independent of the reconstruction procedure. We believe this is the first empirical study to compare different lossy capture strategies without the complication of hardware or reconstruction ambiguity. The scope is limited to linear non-adaptive sensing. The results show that random projections produce significant advantages over other projections only for higher dimensional signals, and suggest more research to nascent adaptive and non-linear projection methods.

1. Introduction

Computational photography involves sophisticated capture methods to capture high dimensional visual signals using invertible multiplexing of signals. This is achieved by careful capture time projection followed by sophisticated

reconstruction. To reduce the number of measurements, a common strategy is to exploit the sparsity in a transformed domain. Let us consider the two projective signal capture approaches for exploiting the sparsity: progressive versus randomized sampling.

Progressive projections Consider the single pixel camera [11]. The successive basis from domains like DCT or Wavelet can be used as modulation patterns to progressively capture higher frequencies. Since signal energy is often compactly represented with the first few coefficients, recovering them is useful for reconstruction. M measurements using these progressive projection patterns lead to recovery of first M frequency coefficients within the basis chosen. The reconstruction is straightforward via weighted combination of each of the M basis. This has been well-known technique for visual signal compression for a long time (JPEG [14], MPEG [18], Transform coding [7], Wavelet based compression [8]).

Random projections Random projection based methods are often studied in the compressive sensing [5][9] and have been applied to various acquisition problems in vision and graphics. Noteworthy among these are: Single pixel camera [11], Compressed video sensing [10], CS Light field capture [3], and Multispectral capture [12]. This has spun a new array of techniques that observe linearly mixed random measurements (projections) and reconstructs the signal using compressive sensing based reconstruction algorithms (L_1 minimization). For signals that can be shown to be sparse in some basis, it has been shown that observing $M = cK \log(\frac{N}{K})$ linear observations is sufficient to ensure exact recovery of such signals [6].

We aim to empirically answer the question: Are random projections based signal capture methods more effective than progressive projections for visual signals? This question, though important has not been answered concretely in the literature due to multiple challenges: (a) The first challenge is to define effectiveness. (b) The second chal-

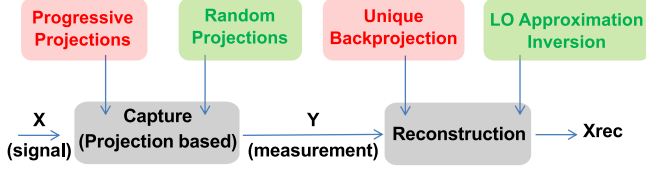


Figure 1. Pipeline for Signal capture and reconstruction using Progressive Transform Coding (TC, Red) or Randomized sparsity coding (SC4, Green).

length is to create a metric that is independent of a specific capture strategy or the sophistication of the (future) reconstruction algorithm. (c) The third challenge is to accommodate nascent adaptive projection methods in comparison. These factors make the discussion on sparsity exploiting methods very tortuous. We address this using the following approach: (a) We compare effectiveness by computing the reconstruction SNR (signal to noise ratio) for a given compression factor. The approach is empirical and we employ large datasets for analysis. Because the datasets are large for analyzing the statistical properties, we also need a fast and efficient method for comparing effectiveness. (b) We use a metric based on Parseval’s energy theorem to relate energy of signal to the energy of the coefficients. Then we use best-case of most recent empirically found recovery bounds to be partial to randomized projection approaches. (c) We limit ourselves to linear and non-adaptive projection methods which are mature and well understood. When the adaptive methods become more established, we agree that a new comparative analysis will need to be performed. In this paper, we design a mechanism to compute the effectiveness for a variety of visual signals such as videos, multi-spectral data and lightfields.

1.1. Motivation

While sparsity has been exploited for image acquisition [11], dual photography [20], reflectance field capture [17] and face recognition [22], there has been little attention paid to the comparison of progressive projections and random projections. The primary analysis in this paper indicates that for images and videos, progressive projection based capture and reconstruction techniques perform comparable to that of random projection based capture and reconstruction. Nevertheless, the analysis also shows that for higher dimensional computational photographic signals such as multispectral data and light-fields, sparsity plays a larger role and techniques that can recover highest magnitude frequency coefficients may be better equipped to tackle these problems, instead of techniques that progressively recover first few coefficients.

Objective and scope: The objective of this paper is to empirically evaluate the relative merits of these sampling techniques. We statistically analyze the compressibility and sparsity structure of visual signals. In an attempt to be invariant to the choice of basis for progres-

sive projection based capture techniques and the choice of basis for sparse signal approximation in random projection based capture, we try several basis functions including DCT, wavelet, Fourier and PCA basis for datasets and compare the best case results. Further, different reconstruction algorithms for compressive sensing basis pursuit, L_1 minimization, matching pursuit methods, etc each have a different performance for different kinds of data. It has been empirically noted that for all of these methods a minimum of $M = 4K$ measurements are required to reconstruct the K -largest magnitude approximation of the sparse signal. This is generally considered an optimistic expectation. In an effort to remain independent of the numerical limitations of the reconstruction algorithms, we do not use any reconstruction algorithm, instead we study the energy of the approximation in $K = M/4$ highest magnitude coefficients, where M is the number of linear measurements. To be favorable to random projections, we:(a) ignore hardware limitations, (b) do not consider effect of quantization. However, the analysis in this paper is limited in the following ways:,we: (a) consider only linear measurements, (d) do not consider structural sparse representations [4], (e) do not consider adaptive methods like- learned dictionaries [15] [19] [21] or overcomplete dictionaries [16] or hybrid sensing techniques [2]

1.2. Contributions

- We devise a procedure to verify effectiveness of progressive as well as random projection based capture methods by comparing the reconstruction SNR for each compression factor. Importantly, the comparison method is fast and independent of capture hardware or reconstruction algorithm.
- We analyze the sparsity, compressibility and energy compaction of several visual signals such as images, videos, multi-spectral imaging and light-fields.
- We empirically demonstrate the analysis using a large dataset for multiple classes of visual signals.

2. Methodology

Traditional signal processing and sampling theory begins with a continuous domain signal $x(t)$ that is then discretized at some sampling rate which is typically greater than the Nyquist rate. In this paper, we consider signals as discrete domain signals represented after they have been sampled and quantized. Further, we only consider a finite time horizon leading to a finite dimensional discrete vector. Let x denote a finite dimensional discrete vector representing the original signal that needs to be compressed.

Projective capture based comparison: Our goal in this paper is to devise a scheme for comparison of signal capture using progressive transform projections that have been

Data Independent		
$S_{trunc}(1:m) = S(1:m)$	$S_{sort} = \text{sort}(S)$ $S_{trunc}(1:m) = S_{sort}(1:m)$	$S_{sort} = \text{sort}(S)$ $S_{trunc}(1:m/4) = S_{sort}(1:m/4)$
$S_{trunc}(m+1:n) = 0$	$S_{trunc}(m+1:n) = 0$	$S_{trunc}(m/4+1:n) = 0$
Transform Coding	Sparse Coding 1	Sparse Coding 4 (CS)

Figure 2. Pseudo codes for computations. The table shows compression techniques as implemented logically. 'S' is the representation of signal in frequency domain. Above pseudo codes signify the possible recoveries of frequency coefficients (S_{trunc}) through: Progressive projection capture (TC: Use of first M coefficients), Random projection capture and L_1 inversion (SC4: Use of highest magnitude M/4 coefficients) and Oracle capture (best possible: SC1: Use of M highest magnitude coefficients) all for $C = \frac{m}{n}$

devised with the ideas of energy compaction in mind. Traditional energy compaction techniques, when used directly in a hardware device observe linear combinations which are directly the dot products with the basis functions. When compression is required they observe far fewer dot products, i.e., dot products only with the first few frequency basis elements (it is well documented that signals usually have more energy in the first few coefficients). The newer emerging techniques of sparse representation and compressive sensing, instead rely on the fact that the signal is sparse in some appropriate domain. In hardware this amounts to observing linear combinations (or dot products) of the signal with some random binary (or Gaussian) vectors. The best estimate of the signal is then obtained via an L1 optimization algorithm. In order to compare the effectiveness of these two techniques for the process of signal acquisition, we devise a scheme that accurately approximates the reconstruction quality of a large class of visual signals. Further, the technique we propose is device and algorithm independent in the sense that limitations of the hardware sensing device or the computational algorithms do not affect our results.

2.1.Choice of Basis

We would like our analyses to be as basis independent as possible. In order to achieve this, we analyze the performance of each of the methods using a variety of commonly used data independent basis such as DCT, wavelet, Fourier etc. Further, in order to evaluate performance characteristics due to the use of data dependent basis, we also use bases obtained via principal component analysis on a separate dataset for each of the datasets on which we evaluate our results. Recently, much progress has been made in the field of dictionary learning, which is another data-dependent basis better suited to sparse representations. While, performance improvements over PCA maybe obtained via the careful use of dictionary learning methods, these methods are yet to reach a stage of maturity that allow easy empirical evaluations. Therefore, in this paper we restrict our attention to data-independent bases and to PCA basis.

2.2.Comparison of projections

For the rest of this section, let us assume that we have a hardware device available with us, capable of observing M measurements, where each measurement is a linear combination of the signal with some known vector. The hardware device is not restricted in the sense that it can obtain a linear combination with any real valued vector. Similarly, we will assume that at the back-end we have computational capacity to reach the optimal solution that is feasible and are not restricted by computational constraints. (Pseudo codes in Figure2).

Progressive Transform Coding (TC): This method attempts to approximate the reconstruction performance that would be obtained while adapting traditional compression techniques to capture devices. In hardware, this would amount to obtaining direct linear combinations with the basis elements (in the order of increasing frequency). Reconstruction, would amount to a backprojection into the signal domain which again is a linear operation. We use the well known Parseval's theorem to approximate the reconstruction fidelity. When M measurements are obtained, this corresponds to the energy captured by the first M frequency basis elements, irrespective of the magnitude of these coefficients. Signals can be represented using their projections onto the basis vectors as $s = [s_1 s_2 s_3 \dots]^T = \Phi x$, where s represents the transform domain representation of the signal. The signal itself can be reconstructed from the basis coefficients as $\hat{x} = \Phi^{-1} s = B s$ since Φ is a full rank matrix. This can also be written as a linear combination of the basis functions as $x = \sum_1^N s_i B_i$, where B_i is the i^{th} basis function. Most common bases in which signals are represented include the discrete cosine transform (DCT), wavelet transform and the Fourier transform. Most of the times, first few components of the representation $s_1, s_2, s_3, \dots, s_K$ contain most of the signal energy. Thus a K-term approximation for the signal can be obtained via $x_{K-term} = \sum_1^K s_i B_i$. This approximation usually captures significant amount of the signal energy. Thus, if each progressive projection pattern were made up of s_i basis vectors, progressively observing 'M' measurements of the progressive projections would be an 'M' term approximation of the signal in respective basis.

Oracle Sparsity Coding 1 (SC1): The goal of this method is to obtain an upper bound for the performance of sparse coding and compressive sensing techniques. As in traditional random projection methods the capture hardware would amount to observing random linear combinations with either random binary or random Gaussian entries. Traditional random projections based reconstruction predicts that if M measurements are observed then one would be able to reconstruct K highest magnitude basis weights where K is given by $M = K \log(N/K)$. Thus, having observed M measurements allows us to only reconstruct $K < M$ basis weights. In the best case scenario, where an

oracle allows us knowledge of the support (i.e. which coefficients have highest magnitude-unavailable in any practical setup), one can reconstruct the M highest magnitude coefficients from the M measurements. Thus the energy captured by M highest magnitude coefficients represents an upper bound on the reconstruction performance using random projections. Most practical methods would perform much poorer than this upper bound.

Randomized SparsityCoding (SC4): As described before, Randomized sparsity coding observes random linear combinations and we expect to reconstruct K highest magnitude coefficients when we solve an L_1 optimization for reconstruction. Here we define that for $s = \phi x$, we measure Y as $y = \psi \phi x$, where S is sparse to some extent and ψ is a random matrix that satisfies the restricted isometry property (RIP). It should be noted that the extent of sparsity in S is not exactly known here (as is the case in most of the visual signals). Here x_{rec} is obtained by solving the linear system $y = \psi \phi x$ using L_1 minimization (basis pursuit). In theory $M = cK \log(\frac{N}{K})$ linear measurements are required where c is an unknown constant. But in practice an efficient and robust algorithm for reconstruction might be able to reconstruct $K = M/4$ highest magnitude coefficients. While, this estimate is also optimistic, it is realistic to expect that current and future algorithmic advances will allow us to reach this. This would amount to the energy captured by the highest magnitude $M/4$ coefficients and serves as a realistic best case performance of Randomized sparsity coding. Thus, this process at its best possible inversion performance is equivalent to Oracle sparsity coding (SC1) but varies in the sense that for a use of M measurements only $M/4$ coefficients effectively contribute. While all other methods explained hitherto are data independent (except the sparsity prior in Randomized sparsity coding), we also experiment with Principal Component Analysis-a data dependent compression technique.

2.3.Evaluation Metrics:

To formally analyze the extent of compression achieved and the quality of reconstructed signals, following statistical parameters can be used. The elements in s are the coefficients for describing the signal in transform domain.

Compression Factor (C): We define the compression factor as the number of coefficients used out of the total number of maximum coefficients. $C = \frac{M}{N}$

Reconstruction SNR (SNR_{rec}): Reconstruction SNR is the Signal to Noise Ratio defined for the reconstruction quality. We define this as the ratio of original signal energy to the error of the reconstructed signal compared to original signal. If original signal is x and reconstructed signal is x_{rec}

$$SNR = 20 \log_{10} \left[\frac{\|x\|}{\|(x-x_{rec})\|} \right] \text{ dB}$$

From Parseval's energy theorem,



Figure 3. Example of compression of an image using competitive compression techniques with a constant compression factor of 0.02. (a) Original image (b) Compression using Progressive projection capture -TC (c) Compression using Random projections capture-SC4 (d) Compression using Oracle Capture-SC1. Performance order is: $SC1 > TC > SC4$. This shows a case where Progressive projection performs better than Random projection.

$$SNR = 20 \log_{10} \left[\frac{\|s\|}{\|s-s_{trunc}\|} \right] \text{ dB}$$

This key idea allows us to compare SNRs directly in sparse coefficient domain without worrying about reconstructing X_{rec} from S_{trunc} . Thus the method is fast and independent of capture hardware or reconstruction algorithm.

3.Datasets

For carrying out an empirical analysis of how random projection capture (compressed sensing) performs in comparison with other compression techniques we do the experiments over competitive methods of signal compression and reconstruction. We elaborate the competitive techniques used for signal compression in the signal processing literature and explain every technique for a generalized case of signal reconstruction (independent of dimensions: images, videos, multispectral data, light fields or signal parameters being compressed: light field spatial resolution, light fields angular resolution). For generality, let us assume that we have total N measurements and for compression we use M measurements ($C = m/n$). Also, let us assume that $s = [s_1 s_2 \dots s_n]$ and $s_{trunc} = [s_{t1} s_{t2} \dots s_{tn}]$. We performed the SNR vs Compression factor analysis of following visual signals over a range of compression factors [0 to 1] for the following datasets: (1) Images (2D) (a) Real images included green channel of 1000 random 2MP Flickr images resized to 256×256 pixels. (b) Cartoon images included green channel of 1000 random 2MP high quality cartoon images resized to 256×256 pixels. (c) Face images included 4596 images from Yale B database [13]. For PCA, the dataset was split into 4096 test images for learning and 500 gallery images for operation.(2) Videos included 30 uncompressed videos resized to $64 \times 64 \times 64$ pixels.(3) Multi Spectral Data included Columbia CAVE labs database, 25 scenes with $512 \text{ pixels} \times 512 \text{ pixels} \times 31 \text{ wavelengths}$. For PCA, cutout versions included 15 by 15 by 31 wavelength (= 6975) data points. (4) Light fields from (New)Stanford light field archive included 14 sets of 17×17 grid (289 views) with $1024 \text{ by } 1024$ resolution (resized to 256×256 pixels). For more details see supplementary material

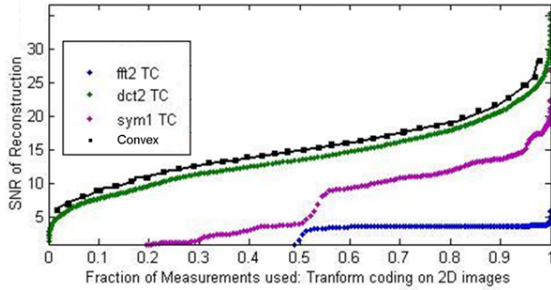


Figure 4. Individual performance plots. We compare the individual performance of wavelets, FFT and DCT for each compression method (TC, SC1, SC4) and plot the best performance curves for each. The black curves show best performance curves.

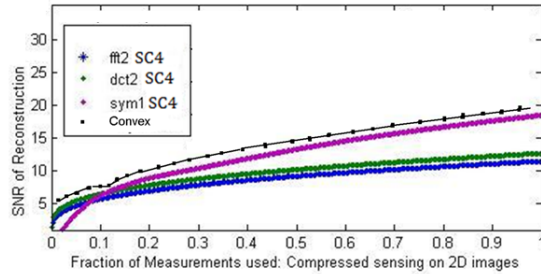


Figure 5. The plot shows how individual best performance curves are used to plot the final result

4. Results

In order to perform the required comparison we analyzed the relative sparsity and energy compaction of visual signals in various data independent basis such as DCT, wavelet, Fourier etc and also the data-dependent PCA basis. We present the reconstruction performance Vs Compression factor plots for: Progressive projection capture (TC), Oracle capture (SC1) and Random projection capture (SC4). For each of the three methods, i.e., TC, SC1 and SC4 we plot the SNR vs compression factor for several different choice of basis. In most cases, the bases that were chosen were DCT, Fourier, Haar Wavelet and Farras Wavelet [1]. Once the separate performance plots for each of the individual bases were obtained, (as shown in Figure 4), then for each compression factor, we only retain the best performing basis. Thus in essence, the performance plots we show are the convex hull (best case) of the performance plots for each of the individual bases as in figure 5.

Interpreting plots: Refer to Figure 6.

4.1. Image datasets

We performed the analysis of energy compaction and sparsity for images using 3 different kinds of image datasets: Random images, cartoons and Faces.

Images: We compared the effectiveness using the following bases: DCT, FFT, Symlet, Daubechies and Haar wavelet. Figure 8 shows the results. Randomized sparse coding method becomes less effective as it does not know

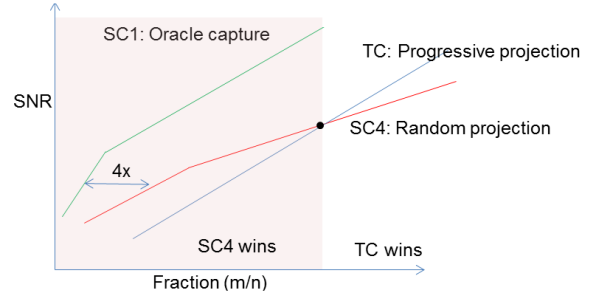


Figure 6. How to interpret plots?: The approach in this paper is to validate Randomized sparsity coding (SC4) method against Progressive transform coding (TC). Horizontal axis shows the compression factor. SC1 is the best-case coding and is stretched in the horizontal direction four times to give SC4. The example here shows a scenario where compressive sensing may work, i.e., SC4 has a 'win' region with respect to TC at highly compressed factors.



Figure 7. Image, Cartoon and face datasets: Image dataset made up of Flickr and internet images, Cartoon dataset: Tom and Jerry, Mickey mouse, Southpark, Face dataset: Yale B face dataset

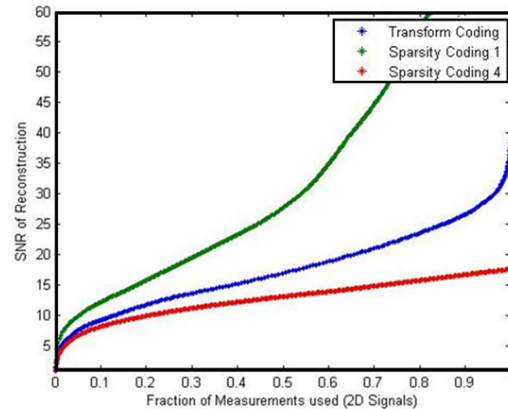


Figure 8. Compression of images gives low reconstruction SNR for lower values of C. Randomized sparsity coding (SC4) performs poorly in comparison with other methods

the non-zero coefficients locations. Progressive transform coding performs better than randomized sparsity coding

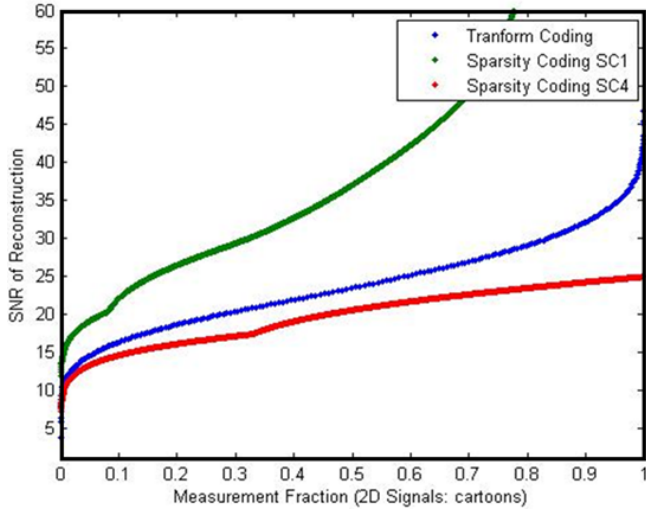


Figure 9. Although cartoons are sparse in gradient domain, TC continues to perform better than SC4.

(SC4) for this dataset.

Cartoon Images: Figure 9 shows the results for experiments on random cartoon image dataset. Cartoons are sparse in gradient domain and have spectral distribution that is very different from real images. This reflects in a better compression performance compared to the images dataset.

Face Images: For the face images, we performed the analysis using the data-independent bases as before (Figure 10). But since faces have a lot of structural similarity we also performed analysis on face images using a data-dependent PCA basis. The basis was learnt from an independent face dataset and there was no overlap between the training and the testing datasets. Data Dependant techniques perform well for face images. Notice that while data-independent random sparse coding (SC4) and data-independent progressive transform coding (TC) seem to perform similarly, the data-dependent basis boosts performance significantly. Nevertheless, there is not a great benefit in performing Randomized sparsity coding (SC4) over Progressive transform coding (TC) even for this dataset.

4.2. Videos

Videos exhibit comparatively more redundancy because of inter-frame overlaps. For videos SC4 does seem to perform marginally better than transform coding at very low compression factors (See Figure 12). Since most interesting compression techniques are concerned with larger compression factors, SC4 does not provide a significant advantage over TC even for such video data. The performance of both these methods maybe improved significantly using motion compensation and other model based methods. The results presented in this paper do not extend directly to such model based methods.

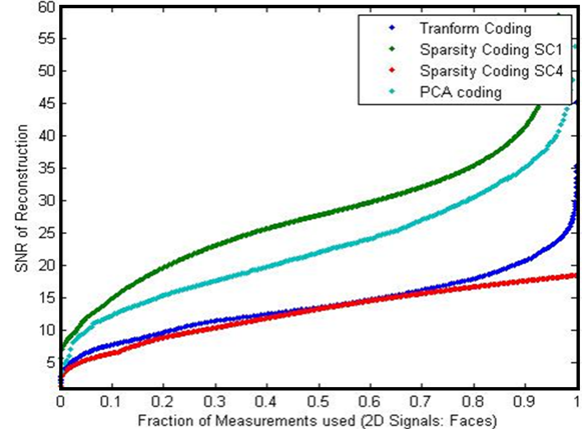


Figure 10. PCA based compression performs well for faces. For other techniques they match with image compression performance




Figure 11. Video dataset- HD documentaries and movies, Multispectral dataset-CAVE labs, Columbia University

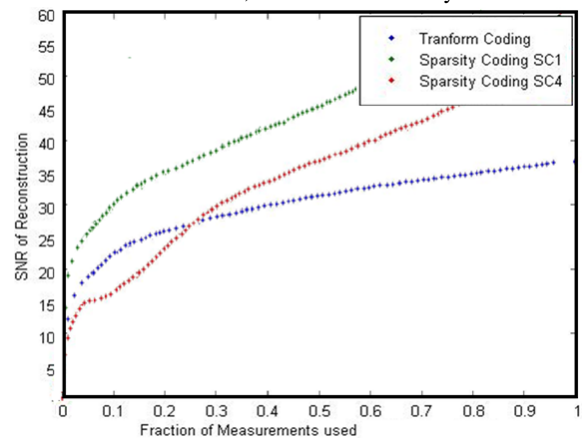


Figure 12. Videos are sparser compared to images, due to inter-frame motion vector redundancies. Compression techniques perform better for videos than for images.

4.3. Multispectral

Figure 13 shows that multi-wavelength capture using Randomized Sparsity coding is beneficial. The intensity

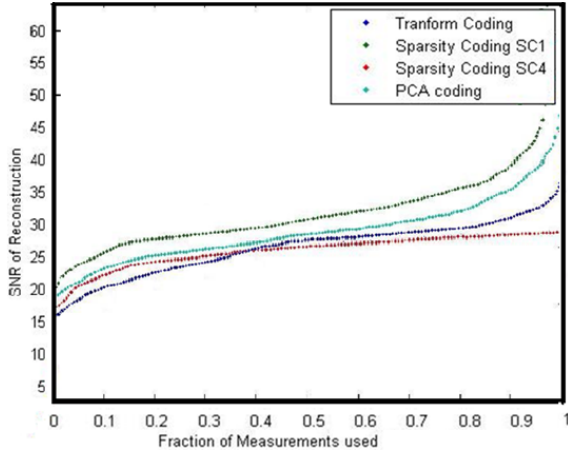


Figure 13. Multispectral data shows intra-frame intensity variations. Compression performance is slightly better for multispectral data compared to images and videos. SC4 starts showing advantages at lower values of C

level variations among frames make the multispectral data sparser in third dimension. Therefore, employing sparsity based compression techniques lead to significant benefits at large compression factors. Since this is the point of operation that is practical, we see that compressive sampling confers a sampling advantage over traditional transform coding for this dataset.

4.4. Lightfields

For the light-fields dataset, we notice in Figure 15 that Randomized Sparsity Coding (SC4) does perform significantly better than Progressive Transform Coding (TC). This indicates that sparse representations and compressive sensing confer a significant benefit over traditional techniques for capture and representation of light-fields.

The information contained in light-fields that is not available in traditional images is subtle disparity information and information about specular highlights. The adjacent views of the light-field are usually very similar and the reconstruction SNR might be high even when this additional information (disparity etc) is not well captured. In order to evaluate whether this information embedded in the light-field is well captured, we also evaluated another metric for light-fields. We analyze the SNR of the center view subtracted Light-Field in order to evaluate the ability of these techniques to capture the angular information. The results are shown in Figure 16. This shows that SC4 performs better than TC even for this angular information metric. Therefore, the disparity information is also better preserved by the Randomized sparsity coding techniques.

5. Conclusions

We chose conditions favorable to projective signal capture, e.g., ignoring capture and reconstruction issues. Noise will be highly amplified after reconstruction. Randomized




Figure 14. Stanford Light field archive (a) Spatial (b) Angular disparity

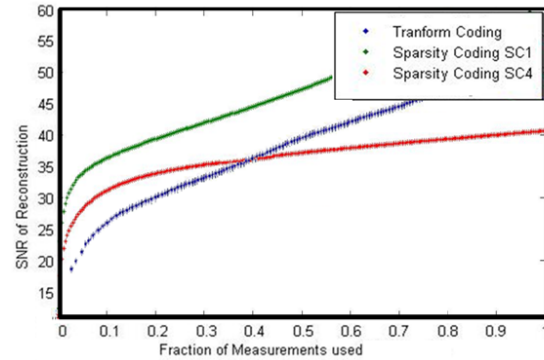


Figure 15. Randomized Sparsity Coding (SC4) shows up to 5 dB better performance compared to Progressive Transform coding (TC). Light field sub-aperture views are sparser by nature and hence allow better reconstruction using SC4

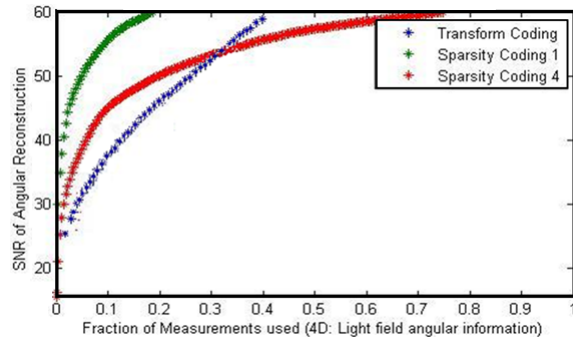


Figure 16. Light fields in angular domain are relatively sparse. The angular sampling of rays does not show significant pixel value variations. SC4 shows notable reconstruction SNR values at lower values of C. The overall reconstruction performance is best among the dataset types tested

projections require basis pursuit algorithms assume near-idealized sensors and hence are highly susceptible to sensor noise or variation from system model. Quantization alone adds one-half of a digital level as noise (For an 8 bit image, the SNR of measured values is clamped at about 53dB.) Randomized projections are more likely to use high frequency patterns in optical path (e.g. coded apertures) making diffraction or calibration more critical. Bayesian inference and other prior based methods may benefit both types of projections. Our discussion was limited to linear and non-adaptive projections. Results of the empirical experi-

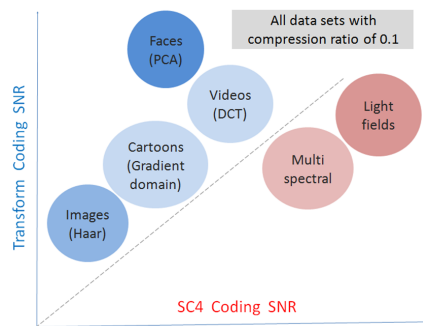


Figure 17. Summary of relative performance plotted as progressive transform coding SNR against randomized sparsity coding SNR for each visual data type. Images have low SNR for both. Higher dimensional signals such as multispectral data and light fields have higher SNRs for SC4 coding.

ments indicate that under idealized sensor and reconstruction conditions, randomized projections offer a benefit for acquiring high dimensional signals such as multi-spectral and light-field data in roughly this order.

images < faces < cartoons < videos < multispectral < LF

Thus more research is needed in novel hardware for randomized projection based approaches for light fields and even reflectance fields.

The sparsity of visual signals is difficult to quantify. By experimenting over large sets, we hope to have captured the characteristics of natural visual signals. There is a remote possibility that the data is in fact significantly more sparse in some unusual yet-to-be-discovered transform basis. But that would be useful for general (non-linear, software based) image compression as well. The lack of such software algorithm indicates low likelihood of such a transform and we consider our empirical results as representative. Hence, researchers should further investigate non-linear or data-adaptive measurements - like use of learned dictionaries or over-complete dictionaries directly suitable for capture purposes.

References

- [1] A. F. Abdelnour and I. W. Selesnick. Symmetric nearly shift-invariant tight frame wavelets. *IEEE Transactions on Signal Processing*, 53(1):231–239, 2005.
- [2] A. Ashok and M. A. Neifeld. Compressive imaging: Hybrid projection design. *Imaging Systems*, page IWD3, 2010.
- [3] S. D. Babacan, R. Ansorge, M. Luessi, R. Molina, and A. K. Katsaggelos. Compressive sensing of light fields. In *IEEE International Conference on Image Processing*, Cairo, Egypt, July 2009.
- [4] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde. Model-based compressive sensing. *CoRR*, abs/0808.3572, 2008.
- [5] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006.
- [6] E. Candes, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *COMMUNICATIONS ON PURE AND APPLIED MATHEMATICS*, 59(8):1207, 2006.
- [7] R. J. Clarke. *Transform Coding of Images*. Academic Press, Inc., Orlando, FL, USA, 1985.
- [8] B. L. B. DeVore, R.A.; Jawerth. Image compression through wavelet transform coding. *IEEE Transactions on Information Theory*, 38(2):719–746, 1992.
- [9] D. L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [10] I. Drori. Compressed video sensing.
- [11] M. F. D. et al. Single-pixel imaging via compressive sampling. *IEEE Signal processing magazine*, 25(2):83–91, 2008.
- [12] M. E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schulz. Single-shot compressive spectral imaging with a dual-disperser architecture. *Opt. Express*, 15(21):14013–14027, 2007.
- [13] A. Georghiadis, P. Belhumeur, and D. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 23(6):643–660, 2001.
- [14] J. Jeong. The jpeg standard. pages 91–99, 1997.
- [15] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online dictionary learning for sparse coding. pages 689–696, 2009.
- [16] J. F. Murray and K. Kreutz-Delgado. Learning sparse over-complete codes for images. *J. VLSI Signal Process. Syst.*, 46(1):1–13, 2007.
- [17] P. Peers, D. Mahajan, B. Lamond, A. Ghosh, W. Matusik, R. Ramamoorthi, and P. E. Debevec. Compressive light transport sensing. *ACM Trans. Graph.*, 28(1), 2009.
- [18] F. Pereira. Mpeg multimedia standards: evolution and future developments. In *MULTIMEDIA '07: Proceedings of the 15th international conference on Multimedia*, pages 8–9, New York, NY, USA, 2007. ACM.
- [19] R. Rubinstein, M. Zibulevsky, and M. Elad. Double sparsity: learning sparse dictionaries for sparse signal approximation. *Trans. Sig. Proc.*, 58(3):1553–1564, 2010.
- [20] P. Sen and S. Darabi. Compressive Dual Photography. *Computer Graphics Forum*, 28(2):609 – 618, 2009.
- [21] Y. Weiss, H. Chang, and W. Freeman. Learning compressed sensing. 2007.
- [22] J. Wright, A. Yang, A. Ganesh, S. Sastry, and Y. Ma. Robust face recognition via sparse representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 210–227, 2009.