

Task Specific Compressive Sensing for Target Detection

Abhijit Mahalanobis

Center for Research in Computer Vision
University of Central Florida
Orlando, FL 32816
UNITED STATES

Mark Neifeld

Department of Electrical and Computer Engineering
University of Arizona
Tucson, AZ 85721
UNITES STATES

ABSTRACT

While the theory of compressive sensing has been very well investigated in the literature, comparatively little attention has been given to the issues that arise when compressive measurements are made in hardware. For instance, compressive measurements are always corrupted by detector noise. Further, the number of photons available is the same whether a conventional image is sensed, or multiple coded measurements are made in the same interval of time. Thus it is essential that the effects of noise and the constraint on the number photons must be taken into account in the analysis, design and implementation of a compressive imager. In this paper, we present a methodology for designing a set of measurement kernels (or masks) that satisfy the photon constraint and are optimum for making measurements that minimize the reconstruction error in the presence of noise. Our approach finds the masks one at a time, by determining the vector which yields the best possible measurement for reducing the reconstruction error. The sub-space represented by the optimized mask is removed from the signal space, and the process is repeated to find the next best measurement. Results of simulations are presented that show that the optimum masks always outperform reconstructions based on traditional feature measurements (such as principle components), and is also better than the conventional image in high noise conditions.

Note: This is an abridged version of the previously published paper “Optimizing Measurements for Feature Specific Compressive Sensing” which can be found at

<https://www.osapublishing.org/ao/abstract.cfm?uri=ao-53-26-6108>

1.0 INTRODUCTION

Measurement is the physical process by which a continuous-valued signal is sampled to yield a number. Any such measurement can be mathematically represented as an inner product between the signal of interest $x(\mathbf{r})$ and a measurement kernel $p(\mathbf{r})$ where \mathbf{r} is a generalized coordinate vector which could include space, time, and wavelength, so that the scalar measurement $u = \int x(\mathbf{r})p(\mathbf{r})d\mathbf{r}$. It is important to note that many/most conventional measurement systems employ a collection of “local” kernels $\{p_i(\mathbf{r}), i=1, \dots, N\}$ each of which has nonzero value only when \mathbf{r} is within some local neighborhood of its “center.” A conventional digital video camera is an example of such a measurement system with each pixel generating a measurement u_i according to a local integral in the transverse spatial dimension (i.e., the lateral resolution), the time dimension (i.e., the frame time), and over some contiguous set of wavelengths (i.e., the color band). It is well-known that images obtained in this way are highly redundant/compressible and for this reason recent work in the domain of compressive sensing (CS) has sought to achieve more efficient use of measurement

resources via use of more arbitrary non-local measurement kernels. So instead of making conventional measurements (i.e., which may be directly consumed by a human observer) CS advocates collecting data in the compressed domain (i.e., in a representation that is not human-viewable) and recovering a more conventional image representation as a post-processing step [1-4]. Although the mathematical underpinnings of CS are very well understood, comparatively little attention has been given to hardware and physical constraints that must be addressed in practice. The CS literature generally uses the same measurement model described above (i.e., after appropriate discretization and signal un-rastering) in which the N -dimensional signal vector \mathbf{x} is projected onto a set of kernel vectors typically arranged as the rows of the $M \times N$ measurement matrix Φ , to yield the M -dimensional measurement vector $\mathbf{u} = \Phi\mathbf{x}$. Unfortunately, much of this theoretical work tends to overlook important constraints associated with the measurement matrices that can be realized in practice and although the theoretical results are impressive, the optimum solutions may not be realizable in practical systems, or yield sub-optimum results when implemented in real-world hardware.

An important constraint associated with any physical measurement system is that the quantity of light energy available to a passive imaging system is limited. In other words, every point in the scene has a fixed brightness (i.e., determined by the scene illumination and reflectance properties at that point) which gives rise to a fixed number of photons being collected by the measurement apparatus within a given collection aperture in space, time, and wavelength. Because this fixed photon budget contributes to several measurements, each such measurement receives only a fraction of the photons from a given point in the scene. Detector noise is also sometimes overlooked in the theoretical CS literature. It is inevitable that measurements contain noise introduced by the detection process which will impact the accuracy of the results. Although previous work has been reported that analyzes the impact of noise on the CS process [5,6], these do not explicitly deal with optimizing the measurement scheme to combat the effects of noise. Neifeld and Shankar [7] were the first to formalize this model, and proposed strategies for feature specific imaging (FSI) using noisy measurements while imposing constraints to meet the photon budget. They observed that direct feature measurement exploits the multiplex advantage, and for small numbers of projections can provide higher feature-fidelity than those systems that post process a conventional image. Since then, further development of the FSI framework has paved the way for compressive imagers that can be realized in optical hardware [6], and information theoretic techniques have been developed for quantifying the information contained in the compressive measurements for performing specific tasks [9]. The primary difference between traditional CS and FSI is in the nature of the sensing matrix employed. In particular, CS assumes sparsity in an unknown basis, and the measurements are designed to be incoherent [10]. In FSI, on the other hand, the sensing matrix is designed using all available prior statistical knowledge of the scene. Either paradigm can be appropriate for a given situation, but if the particular exploitation task and the prior scene statistics are known, then the FSI paradigm can be very useful.

Motivated by the FSI framework, we revisit the task of image reconstruction from compressive measurements informed by signal priors. We begin by observing that given a set of noisy measurements, the optimum linear operator to minimize the reconstruction mean square error (MSE) is well known [11]. It is also important to observe that when second order signal statistics are available, the principle component analysis (PCA) features are known to be optimal for image reconstruction in a minimum mean square error (MMSE) sense [12]. Thus, it may be tempting to make measurements using PCA projections and then using the LMMSE operator to reconstruct images; however, this approach would be sub-optimal even in the LMMSE sense. The key insight here is that the derivation of the PCA projections does not take into account the effect of measurement noise. In fact, it has been shown previously [9] that the PCA is indeed sub-optimal under high noise conditions, and that the optimum measurement kernels are binary-valued. In this paper, our goal is to systematically determine the projections which are optimum for making measurements such that the effect of noise on the reconstructed imagery is minimized, subject to the finite photon budget.

The rest of the paper is organized as follows. In Section 2, we present a strategy for mask optimization that takes noise and the photon constraint into account. Assuming a linear MMSE reconstruction model and prior

information based on second order statistics, we show that the optimum masks are binary-valued in high noise conditions, but converge to the PCA vectors when noise is absent. For a given signal to noise ratio (SNR), the algorithm sequentially determines one optimum mask at a time, removes its contribution from the signal space, and then iterates to find additional masks until the reconstruction error is acceptably small. The results of simulation are presented in Section 3. We find that when SNR is low, the optimized masks not only outperform the PCA, but also yield better results than a conventional imager. Under these conditions, FSI with optimum masks yields a smaller MSE than either the reconstruction obtained with PCA masks or a conventional noisy image at comparable SNR. In high SNR conditions, the reconstruction accuracy of the optimum masks is always equal or better than that of the PCA for any given number of measurements. Thus, in addition to ensuring that the photon constraint is satisfied, the optimization technique described in Section 2 enables the masks to achieve robust performance in noise without adversely affecting the reconstruction MSE. Although our analysis assumes that the noise is independent of the signal, this is not the case when Poisson noise is present. While a formal treatment of the effects of signal dependent noise is beyond the scope of the paper, we compare the behavior of optimized masks perform to the normalized PCA in signal dependent noise. Section 4 is a summary of the findings of the paper, and suggestions for direction for future work.

2.0 SYSTEM DESCRIPTION AND MASK OPTIMIZATION

Before discussing the procedure for optimizing the measurement kernels, we briefly review the architecture of a single pixel compressive sensor [4] so that its implications on the optimization framework are understood. Assume that we wish to compute the quantity $u = \sum_{m,n} x(m,n)\phi(m,n)$, where $x(m,n)$ is an image of a scene, and $\phi(m,n)$ is the measurement kernel (interchangeably referred to as a *mask*), and m and n are the discrete spatial indices. As shown in Figure 1, this operation can be optically implemented by imaging the scene on a spatial light modulator, and collecting the modulated light on a photo-detector. The SLM is either a transmissive or reflective device which is encoded with the mask pattern $\phi(m,n)$. The light emerging from the SLM is proportional to the product $x(m,n)\phi(m,n)$, and is collected by the integrating lens on the photo-detector. Effectively, the current measured at the output of the photo-detector is proportional to u .

Since real world SLMs can only implement non-negative masks, the values of $\phi(m,n)$ must range between 0 and 1, and represent the fraction of the light that is transmitted by the mask at each spatial location. On the other hand, the performance of a mask that contain both positive and negative values may be better than those which are strictly positive. To handle the negative values, the mask can be expressed in terms of two non-negative quantities, i.e,

$$\phi(m,n) = \phi^+(m,n) - \phi^-(m,n),$$

where

$$\phi^+(m,n) = \begin{cases} \phi(m,n) & \text{if } \phi(m,n) > 0 \\ 0 & \text{otherwise} \end{cases}$$

and

$$\phi^-(m,n) = \begin{cases} |\phi(m,n)| & \text{if } \phi(m,n) < 0 \\ 0 & \text{otherwise} \end{cases}$$

and the final result can be obtained by subtracting a partial measurement made using $\phi^-(m,n)$ from another obtained using $\phi^+(m,n)$. The optimization technique described in this paper can be formulated to

yield both non-negative or real-valued masks.

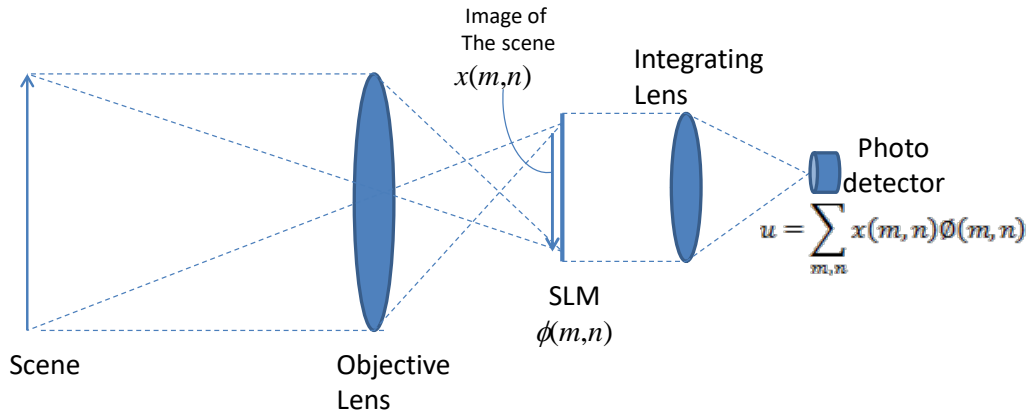


Figure 1: Basic architecture of a single pixel compressive sensor which integrates the light from the scene, modulated by codes written on the SLM

The noise characteristics of a photo-detector depends on the type of device and the wavelength of the light. In particular, we are interested in detectors that operate in the mid-wave infra-red. Cooled sensors are much more sensitive and exhibit relatively low levels of noise that has predominantly signal dependent Poisson characteristics. Thus cooled detectors provide significantly higher signal to noise ratios in many applications, and the impact of noise on CS measurements is less severe. On the other hand, uncooled sensors (such as Lead Selenide detectors [13]) are relatively inexpensive, but are also substantially more noisy. In such cases, the dominant noise is *pink noise* (also known as $1/f$ noise) which does not depend on the signal, and can be modeled as an additive process. The analysis presented in this paper deals with impact of such signal independent additive noise on compressive measurements, and is motivated by the goal of using inexpensive infra-red detectors in compressive sensors.

2.1 Mathematical Framework

For the sake of brevity, we omit the detailed mathematical formulation and treatment of the problem which can be found in the published version at “<https://www.osapublishing.org/ao/abstract.cfm?uri=ao-53-26-6108>”

3.0 RESULTS OF SIMULATION

A database of thermal images of six different types of vehicular targets, some of which are shown in Figure 2, was used to create the noise optimized masks for feature specific imaging. It should be noted that such thermal images in the long wave infra red spectrum have relatively low spatial frequency content, and the information contained in the pixels is often redundant. Our goal is to examine the results of FSI using far fewer measurement than the number of conventional pixels contained in these image. The images are of size 20×40 pixels, and consequently the target correlation matrix \mathbf{R}_x estimated from these images is of size 800×800 . Since the location of the target in the scene is assumed to be unknown, all possible shifted versions the targets (within the 20×40 window) was also included in the estimation of \mathbf{R}_x .




Figure 2: Typical images used for estimating target statistics

To compare the results of compressive sensing and reconstruction to that of a conventional imaging sensor, we define a SNR in terms of the signal and the detector noise in a conventional camera. Assume that the ideal image x obtained using a conventional imager is corrupted by detector noise that is zero mean AWGN with standard deviation σ . If x_{max} and x_{min} represent the maximum and minimum signal amplitudes, then the signal to noise ratio (SNR) is defined as $(x_{max} - x_{min})/2\sigma$. Once σ is determined, the same noise level is used for making the compressive measurements, and for assessing the performance of the masks. For illustrative purposes, a conventional image of a target with no noise is shown in Figure 3, along with versions of the same image corrupted by noise at SNR values of 5, 30, and 90. For the purposes of our analysis, these are considered to be the high, moderate, and low noise cases, respectively.




Figure 3: Conventional images of a target for various noise levels: a) no noise, b) SNR=5, c) SNR=30, and d) SNR = 90.

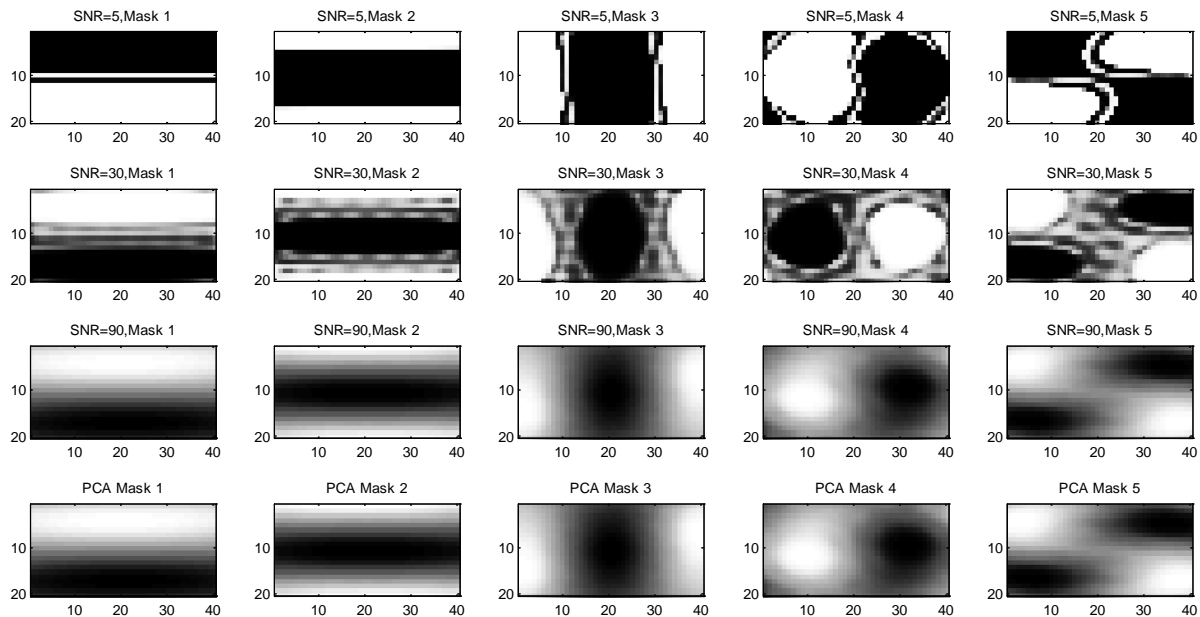


Figure 4: The masks in the top three rows are optimized for SNR values of 5, 30, and 90 respectively. The PCA masks are shown in the bottom row. Note that the masks appear to be almost binary valued for SNR=5 (high noise) and essentially the same as the PCA for SNR=90 (low noise).

We created three sets of masks optimized for various SNR values of 5, 30 and 90 using the algorithm described in Section 2. The first five masks for each case are shown in Figure 4, along with the masks based on the PCA vectors. It is interesting to note that for the low SNR of 5 (high noise case), the optimum masks are practically binary valued. At the moderate SNR of 30, the masks are more saturated than the PCA, but also exhibit some fractional gray values. However, at high SNR of 90 (low noise case), the masks appear to be essentially the same as the PCA vectors.

The logarithm of the reconstruction MSE obtained using these masks for the high noise case is shown in Figure 5. Here, σ was chosen to yield an SNR of 5, and the MSE in Eq. (5) is plotted for different numbers of measurements. In all cases, the MSE is reduced as the number of measurements is increased. It is clear that the masks optimized to handle the high noise case perform the best and always yield the smallest MSE. On the other hand, the normalized PCA yields the worst performances, as expected. It is also noteworthy that the masks optimized for moderate and low noise cases also outperform the normalized PCA, although to a lesser extent. Hence, even if the masks are optimized for SNR values that are different from the actual noise level, their performance is better than that of the normalized PCA when the measurements are corrupted by noise.




Figure 5: In high noise, the masks optimized for an SNR of 5 yield the smallest MSE. Masks optimized for moderate and high SNR also perform better than the PCA.

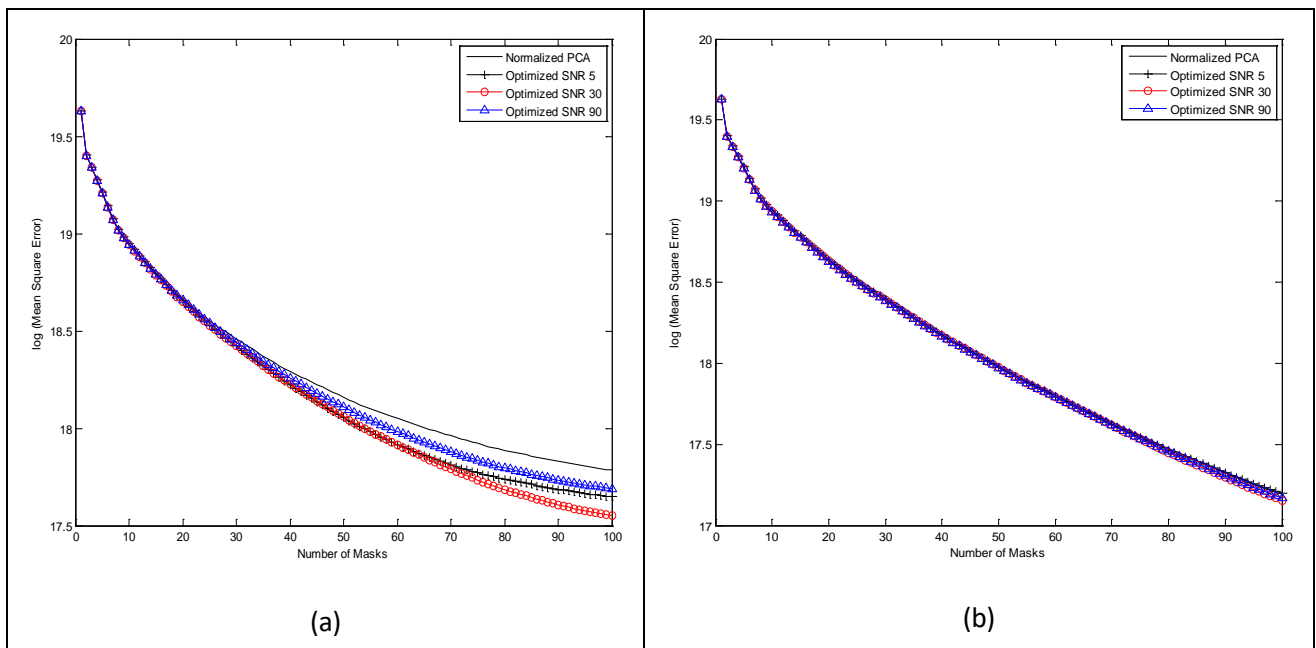


Figure 6: In moderate noise (a), the mask optimized for SNR=30 yields the best performance as the number of features is increased, while the PCA yields highest MSE. For SNR = 90 (b), (i.e. in low noise) the performance of all sets of masks is comparable.

The performance of the masks in moderate and low noise levels are shown in Figures 6(a) and 6(b), respectively. In Figure 6(a), σ is chosen to correspond to an SNR of 30. Although all sets perform comparably with the few initial measurements, we see that as the number of features increases, the set optimized for moderate noise level (i.e. for SNR=30) outperforms the others. Not surprisingly, the normalized PCA produces the largest MSE as the number of measurements is increased. In low noise conditions, the performance of all masks is comparable to the PCA, as shown in Figure 6(b). This implies that the optimization algorithm does not produce any appreciable loss of performance in the absence of noise, but yet provides considerable improvements when noise is present in the measurements.

We now discuss the effect of imposing a photon constraint on the measurements made using the optimized masks. In infra-red (IR) systems, the *integration time* of the detector determines the number of photons collected over that time interval. The plots shown in Figure 5 and 6 assume that the same integration time is used for every measurement. This implies that as the number of masks increases, more photons are utilized in the final reconstruction result. Whether this is feasible depends on the design of the imaging system. For example, if a final frame rate of no more than 10 Hz is desired using an integration time of a millisecond per measurement, then up to 100 masks can be used for reconstructing each image frame. In this scenario, images reconstructed with fewer measurements essentially receive less photons, and the time required for making all the measurements depends on the number of masks to be used.

On the other hand, if all measurements must be made within a fixed interval of time, then the integration time for each measurement depends on the number of masks to be employed. Essentially, this ensures that the number of photons utilized for reconstructing a complete image is always the same. The effect of reducing the integration time for each mask can be treated as a multiplicative scale factor applied to the corresponding measurement prior to adding noise. Therefore, if N measurements are used to reconstruct an image, noise with the same standard deviation σ (determined as before based on the SNR of a conventional image) is added to each measurement scaled by $1/N$. Figure 7 shows the effect of imposing this photon constraint on the behavior of the reconstruction MSE for the high noise case (SNR=5). Under these circumstances, we see that the MSE initially decreases as more masks are used until a minimum value is achieved for 17 measurements. This occurs as long as additional measurements contain more signal information than noise. After that, the MSE starts to increase when more measurements are included since these introduce more noise than useful signal information. It is noteworthy that in Figure 7, the masks optimized for high noise (i.e. the plot indicated with the '+' symbol) outperforms the PCA (solid line) and other masks. This set not only provides a smaller MSE at any number of measurements, but also achieves the smallest overall MSE, and hence provides the best possible reconstruction. The results for the moderate and low noise cases are shown in Figure 8. Here, the mask optimized for SNR=30 yields the smallest reconstruction MSE in moderate noise using 85 measurements as shown in Figure 8(a). For SNR=90 there is no appreciable difference between the performance of the masks as shown in Figure 8(b). After comparing these results to those obtained earlier in Figures 5 and 6, we conclude that although the photon constraint impacts the behavior of the MSE as the number of measurements is increased, the best results for a given SNR is obtained using the masks that were optimized for that noise level. It should be noted that if two measurements are made for every mask (to account for +/- values) based on architectural considerations, then the integration time for each measurement is reduced by a factor of 2, which in turn reduces the measurement SNR. Although this results in a larger MSE for all masks, the overall trends shown in Figures 7 and 8, and our conclusions remain the same.




Figure 7. When a photon constraint is imposed by limiting the integration time allocated to each mask, the MSE initially decreases but then increases again as more noisy measurements are included. In high noise conditions, the best results are obtained using the masks optimized for SNR=5, although the mask optimized for low and medium SNR still outperform the PCA.

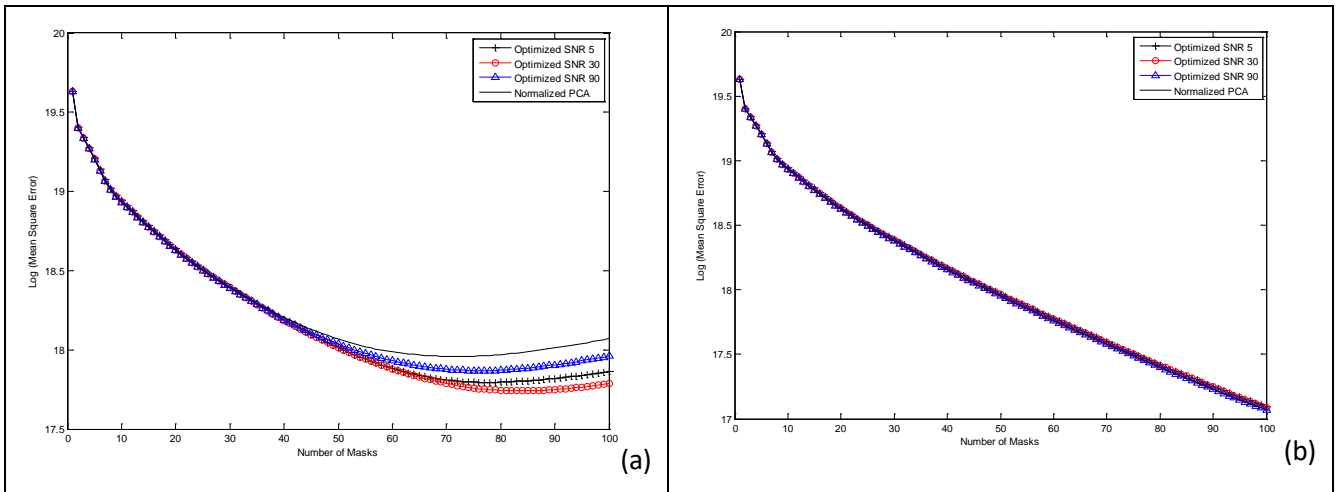


Figure 8: Under the photon constraint at an SNR =30, the mask optimized for moderate noise yields the best result, compared to the PCA and the masks optimized for other noise levels, as shown in (a). In low noise conditions (SNR=90), there is no appreciable difference between any of the masks, as shown in (b).

Although the masks have been optimized for signal independent noise, we now examine their behavior in noise that has a signal dependent component. Towards this end, Eq (1) is modified so that the measurement obtained with the mask Φ is given by $u = \Phi^T x + v + \eta$, where the term η represents a zero mean signal dependent noise with standard deviation $\sigma_\eta = \gamma\sqrt{|\Phi^T x|}$, and v is the signal independent noise process (as before) with standard deviation σ . Thus, σ_η is proportional to the square root of the magnitude of the

projection of the signal on the mask. We chose the constant γ such that $\sigma_\eta = \sigma$ when Φ is the first mask, and is left unchanged for all other masks.

The experiment described in Figure 7 is repeated using the modified noise model, and the results are shown in Figure 9. While the minimum MSE is somewhat greater for all masks, their behavior differs significantly as the number of measurements is increased. The masks optimized for higher levels of noise exhibit a relatively smaller growth in MSE as the number of measurements is increased. For instance in Figure 9, the normalized PCA achieves the minimum $\log(\text{MSE})$ of 19.25 for about ten measurements, whereas the mask optimized for SNR=5 achieves an even smaller minimum $\log(\text{MSE})$ of 19.1 with about 13 measurements. It is clear from these results that the masks designed to minimize MSE in signal independent noise continue to perform better than the normalized PCA, even when signal dependent noise is present.




Figure 9: The masks designed to minimize MSE in signal independent noise continue to perform better than the normalized PCA, even when signal dependent noise is present.

For visual comparison, several examples of reconstructed and conventional images are shown in the table in Figure 10 for different noise levels. For each case, we also note the MSE with respect to the noise-free ideal image, as well as the effective *compression ratio* (i.e. the ratio between the number of pixels in the image, and the number of measurements that were used to obtain the reconstruction). Under very noisy conditions at SNR=5, the optimized masks yields the smallest MSE (20.15) outperforming both the PCA (MSE=28.8) as well as a conventional imager (MSE=26.2). Additionally, only 17 measurements were used to obtain this result, which represents a compression ratio of 47. The reconstruction using the optimized masks has the best visual quality; the reconstruction using the PCA does not resemble the original object, and the conventional image appears to be very noisy.

In moderate noise (SNR=30), the optimized masks (MSE=11.1) still outperform the PCA (MSE=14.6). Although the conventional image has a lower MSE, the reconstruction using the optimized mask and the PCA provide a compression ratio of 9.4. The shape of the target obtained with the optimized masks is visually comparable to the ideal image, whereas the PCA based reconstruction appears more distorted. The target can be clearly seen in the traditional image albeit the noise is also visually evident. Finally, in low noise conditions, very comparable results are obtained using both the optimized masks (MSE=4.61) and the PCA masks (MSE=4.62). Both exhibit a slightly larger MSE than the conventional image (MSE=1.5), mainly because of the compressive nature of the measurements (compression ratio=8).

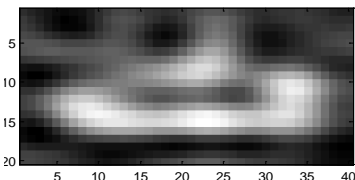
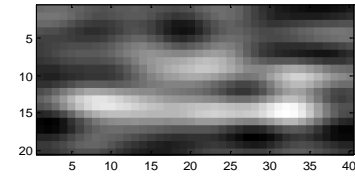
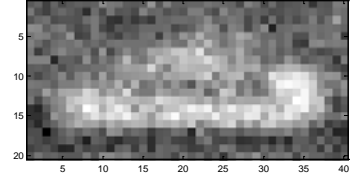
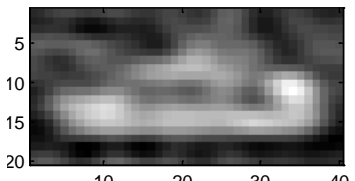
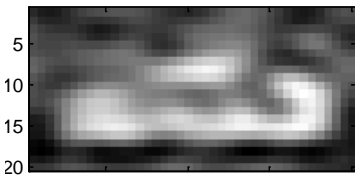
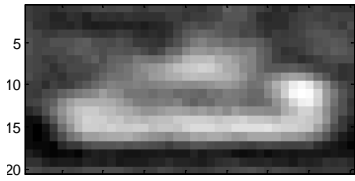
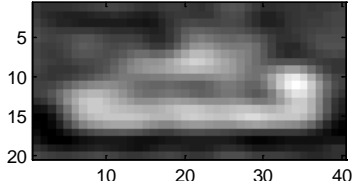
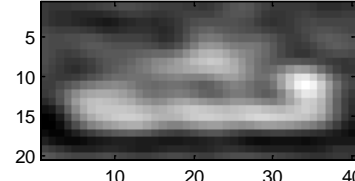
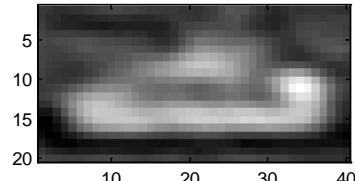
| SN R | Optimized Mask | Normalized PCA | Conventional Image (no compression) |
|------|--|---|--|
| 5 |  <p>MSE=20.15 Compression Ratio= 47</p> |  <p>MSE=28.8 Compression Ratio =47</p> |  <p>MSE=26.2</p> |
| 30 |  <p>MSE=11.1 Compression Ratio = 9.4</p> |  <p>MSE=14.6 Compression Ratio=9.4</p> |  <p>MSE=4.6</p> |
| 90 |  <p>MSE=4.61 Compression Ratio=8x</p> |  <p>MSE=4.62 Compression Ratio=8x</p> |  <p>MSE=1.5</p> |

Figure 10: The results of reconstructing the ideal image in Figure 3(a) using noisy feature specific measurements (using the optimized mask and masks based on the PCA) are compared to the conventional image at the same SNR. The results show that the optimized masks always outperform the PCA by yielding a smaller MSE at the same compression ratio. The results are also better than the conventional image in high noise, and visually comparable to the conventional image in moderate and low noise conditions. In these cases, reconstruction based on FSI exhibit a residual MSE due to the compressive nature of the measurements.

4.0 SUMMARY

While the theory of compressive sensing has been very well investigated in the literature, comparatively little attention has been given to the issues that arise when compressive measurements are made in hardware. For instance, compressive measurements are corrupted by detector noise. Further, the number of photons available is the same whether a conventional image is sensed, or multiple coded measurements are made in the same interval of time. Thus it is essential that the effects of noise and the constraint on the number

photons must be taken into account in the analysis, design and implementation of a compressive imager.

Feature specific imaging (FSI) is a form of compressive sensing where the measurement kernels are not random, but are based on prior knowledge of the information we are interested in sensing. Working in this FSI framework, we have developed a methodology for designing a set of masks that satisfy the photon constraint and are optimum for making measurements that minimize the reconstruction MSE in the presence of noise. The measurement matrix is obtained by arranging these masks as its rows. A criteria for meeting the photon constraint is that the maximum L1 norm of the columns of the measurement matrix is less than or equal to 1.0. To simplify the optimization process, we employed an analytical mapping that ensures the masks can take on any value between ± 1 and formulated a quadratic objective function that can be minimized using gradient descent. The process then finds the mask one at a time, by determining the vector which yields the best possible measurement for reducing the MSE. The sub-space represented by the optimized mask is removed from the signal space, and the process is repeated to find the next best measurement.

Of course, the primary reason for compressive sensing is to make fewer measurements than the number of pixels in the reconstructed image, and to collect the information more efficiently than a conventional image. One might intuitively expect that utilizing more features (or measurements) in the reconstruction process will yield a smaller reconstruction error. We demonstrated however that the photon constraint limits the number of masks that can be used at a particular SNR to reduce the reconstruction MSE. In noisy conditions, MSE initially decreases as the number of measurements is increased, but then increases when measurements that contain more noise than signal information are included. The simulations showed that the optimum masks always outperforms masks based on the PCA vectors, and also yield a smaller MSE than the conventional image in high noise conditions. Although a formal treatment of the effects of signal dependent noise is beyond the scope of the paper, we found that the optimized masks perform better than the normalized PCA, even in signal dependent noise.

Future work will focus on extending the optimization algorithm so that multiple masks are jointly optimized subject to the photon constraint, rather than one at a time. Other types of noise models will be also considered such as signal dependent shot noise. This material is based upon work supported by DARPA and the SPAWAR System Center Pacific under Contract No. N66001-11-C-4092. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

5.0 REFERENCES

- [1] Emmanuel Candès and Terence Tao, "Near optimal signal recovery from random projections: Universal encoding strategies", IEEE Trans. on Information Theory, 52(12), pp. 5406 - 5425, December 2006)
- [2] David Donoho, "Compressed sensing", IEEE Trans. on Information Theory, 52(4), pp. 1289 - 1306, April 2006
- [3] Richard Baraniuk, "Compressive sensing", IEEE Signal Processing Magazine, 24(4), pp. 118-121, July 2007
- [4] Michael Wakin, Jason Laska, Marco Duarte, Dror Baron, Shriram Sarvotham, Dharmpal Takhar, Kevin Kelly, and Richard Baraniuk, "An architecture for compressive imaging", Int. Conf. on Image Processing (ICIP), Atlanta, Georgia, October 2006

- [5] P. T. Boufounos, M. F. Duarte, and R. G. Baraniuk, "Sparse Signal Reconstruction from Noisy Compressive Measurements Using Cross Validation," Proceedings of the IEEE Workshop on Statistical Signal Processing (SSP), pages 299-303, August 2007
- [6] M. Raginsky, R.M. Willett, Z.T. Harmany, R.F.Marcia, "Compressed Sensing Performance Bounds Under Poisson Noise," IEEE Transactions on Signal Processing, Volume 58 Issue 8, pages 3990 - 4002, August 2010
- [7] Mark A. Neifeld and Premchandra Shankar, "Feature-Specific Imaging", Applied Optics, Vol. 42, Issue 17, pp. 3379-3389 (2003)
- [8] Dharmpal Takhar, Jason Laska, Michael Wakin, Marco Duarte, Dror Baron, Shriram Sarvotham, Kevin Kelly, and Richard Baraniuk, "A new compressive imaging camera architecture using optical-domain compression", Computational Imaging IV at SPIE Electronic Imaging, San Jose, California, January 2006
- [9] Amit Ashok, Liang-Chih Huang, and Mark A. Neifeld, "Information optimal compressive sensing: static measurement design", JOSA A, Vol. 30, Issue 5, pp. 831-853 (2013)
- [10] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin, "A simple proof of the restricted isometry property for random matrices", Constructive Approximation, 28(3), pp. 253-263, December 2008
- [11] H.C Andrews, B. R. Hunt, *Digital Image Restoration*, Prentice-Hall signal processing series, 1977
- [12] Clarke, R.J., *Transform Coding of Images*, Academic Press, San Diego, CA, 1985
- [13] G. Vergara, G. VPD PbSe Technology fills the existing gap in uncooled, low cost and fast IR imagers. Proc. of SPIE vol 8012, p.146, (2011).

