ELEG 867 - Compressive Sensing and Sparse Signal Representations

Gonzalo R. Arce
Depart. of Electrical and Computer Engineering
University of Delaware

Fall 2011

Outline

- Applications in CS
- Single Pixel Camera
- Compressive Spectral Imaging
- Random Convolution Imaging
- Random Demodulator

Imaging as the Origins of CS

Magnetic Resonance Imaging

- MRI measures frequency domain image samples
- Fourier coefficients are sparse
- Inverse Fourier transform produces MRI image
- Time of acquisition is a key problem in MRI

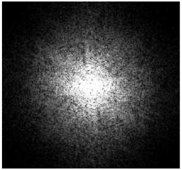
Coefficients in Frequency

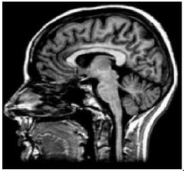
MRI Image

M. Lustig, D. Donoho and J. M. Pauly. Sparse MRI: the application of compressive sensing for rapid MRI imaging Magnetic Resonance in Medicine. Vol. 58. 1182-1195. 2007.

MRI Reconstruction

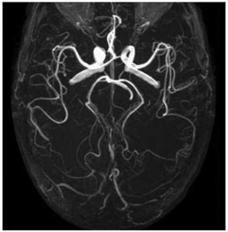

Space

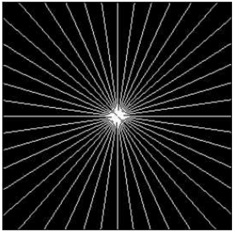

Frequency

- Want to speed up MRI by sampling less. In a N by N image
  - 22 radial lines
  - N Fourier samples for each line
  - If $N = 1024$, 98% of the Fourier coefficients are not sampled

Reconstruction Example

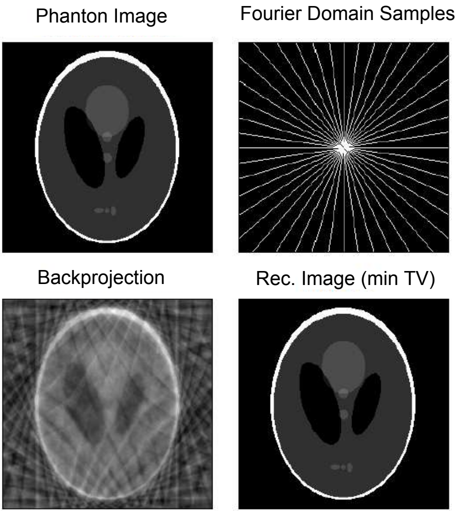

MRI Reconstruction: Formulation Problem

- Reconstruction by minimization of total variation *(min-TV)* with quadratic constraints †

$$\min_x \|x\|_{TV} \ \text{ s.t. } \ \|\Phi x - y\|_2^2 \leq \epsilon$$

  - x is the unknown image
  - $\Phi = F_p$, is the partial Fourier matrix
  - y is the partial Fourier coefficients
  - $\|x\|_{TV} = \sum_{i,j} |\nabla x(i,j)|$ where
    $|\nabla x(i,j)|$ is the Euclidean norm of $\nabla x(i,j)$
- The total variation of the image x ($\|x\|_{TV}$) is the sum of the magnitudes of the gradient.

† E. Candès, J. Romberg and T. Tao "Stable Signal Recovery from Incomplete and Inaccurate Measurements." Comm. on Pure and App. Math. Vol.59,No.8, 2006.

Single Pixel Camera †

- Obtain an image by a single photo detector.

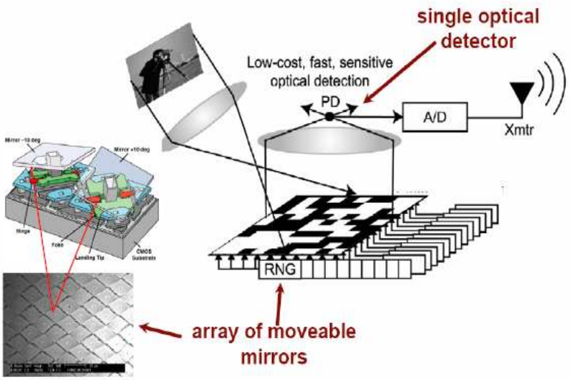

† M. Duarte, M. Davenport, D. Takhar, J. Laska, T. Sun, K. Kelly, and R. Baraniuk. "Single-Pixel Imaging via Compressive Sampling." IEEE Signal Processing Magazine. 2008.

Single Pixel Camera at UD Lab.

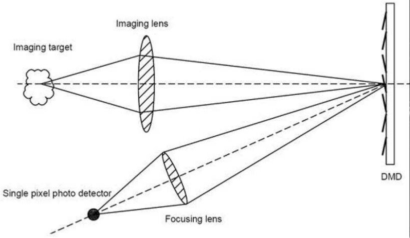

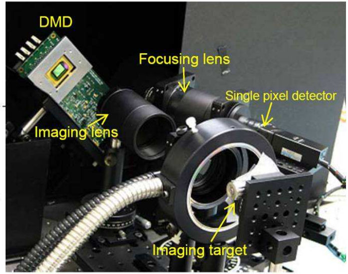

- Incident light field (corresponding to the desired image) is reflected off a digital micro-mirror device (DMD) array.
- The mirror orientations are defined by the entry of the modulation patterns (B^k).
- Each different mirror pattern produces a voltage at the single photodiode (PD) that corresponds to one measurement.

Single Pixel Camera at UD Lab.

3 by 4 mirror sub-arrays

2 by 2 mirror sub-arrays

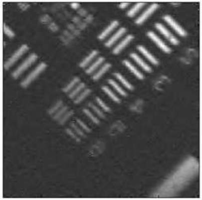

1 by 1 mirror sub-arrays

Single Pixel Camera at UD Lab.

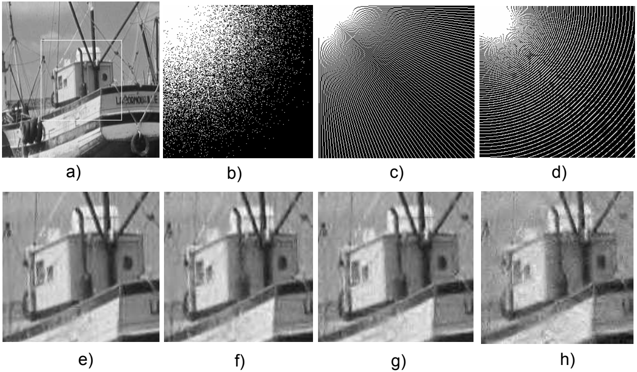

(a) Original, Sampling with (b) Variable density, (c) Radial, (d) Log. spiral. All 30.5% undersampling ratio. Reconstruction with (e) variable density, (f) radial, (g) log. spiral (h) SBHE.

Z. Wang et al. Variable Density Compressed Image Sampling. IEEE Trans. Image Processing, vol. 19, no. 1, Jan.2010.

Compressive Spectral Imaging

- Collects spatial information from across the electromagnetic spectrum.
- Applications, include wide-area airborne surveillance, remote sensing, and tissue spectroscopy in medicine.

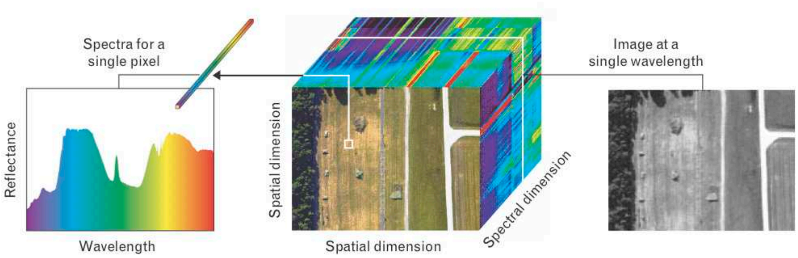

Hyper-Spectral Imaging (HSI)

HSI systems collect information as a set of images. Each image represents a range of the spectral bands. Images are combined in a three dimensional hyperspectral data cube. Scanning HSI sensors use linear detector arrays and a mirror that scans in the cross-track direction to acquire a 2D multi-band image. The linear detector array records the spectrum of each ground resolution cell.

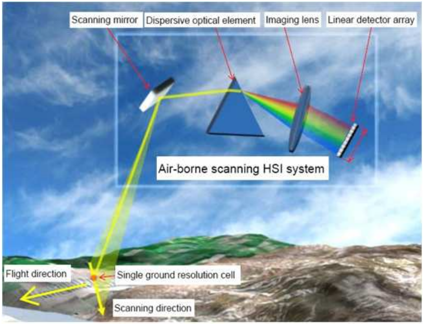

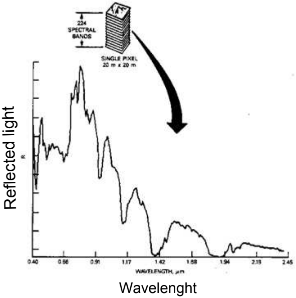

Pushbroom HSI sensors

A 2D array detector is used so that the spectral information of the entire swath width can be collected simultaneously. It does not need moving parts for air-borne or space-borne HSI applications and it has longer dwell time and improved SNR performance.

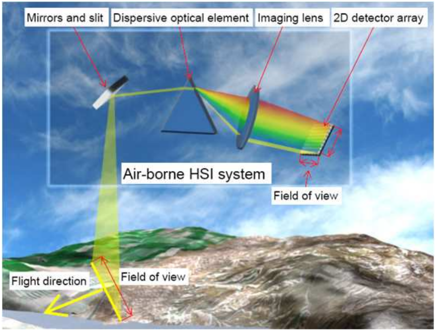

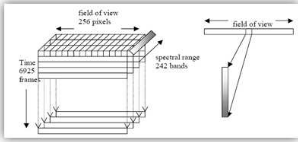

Datacube of the HSI system

Spectral Imaging System - Duke University†

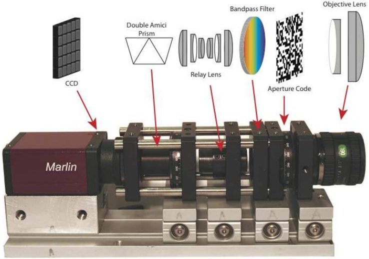

Schematic and photo showing the entire CASSI instrument. Left to right: CCD, Double Amici Prism, Relay Lens, Bandpass Filter, Aperture Code, and Objective Lens

A. Wagadarikar, R. John, R. Willett, D. Brady. "Single Disperser Design for Coded Aperture Snapshot Spectral Imaging." Applied Optics, vol.47, No.10, 2008.

A. Wagadarikar and N. P. Pitsianis and X. Sun and D. J. Brady. "Video rate spectral imaging using a coded aperture snapshot spectral imager." Opt. Express, 2009.

Single Shot Compressive Spectral Imaging

System design

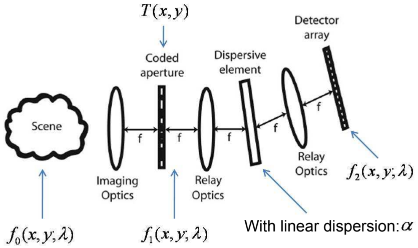

$$
\begin{aligned}
f_1(x,y;\lambda) &= f_0(x,y;\lambda)T(x,y) \\
f_2(x,y;\lambda) &= \int\int \delta(x' - [x+\alpha(\lambda-\lambda_c)]\delta(y'-y)f_1(x',y';\lambda))dx'dy' \\
&= \int\int \delta(x' - [x+\alpha(\lambda-\lambda_c)]\delta(y'-y)f_0(x',y';\lambda)T(x,y))dx'dy' \\
&= f_0(x+\alpha(\lambda-\lambda_c),y;\lambda)T(x+\alpha(\lambda-\lambda_c),y)
\end{aligned}
$$

Single Shot Compressive Spectral Imaging

Experimental results from Duke University

Original Image

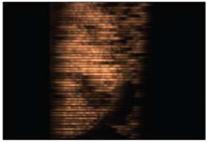

Measurements

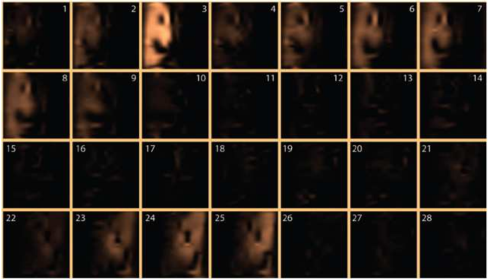

Reconstructed image cube of size:128x128x128.
Spatial content of the scene in each of 28 spectral channels between 540 and 640nm.

† A. Wagadarikar, R. John, R. Willett, D. Brady. ”Single Disperser Design for Coded Aperture Snapshot Spectral Imaging.” Applied Optics, vol.47, No.10, 2008.

Single Shot Compressive Spectral Imaging

Simulation results in RGB

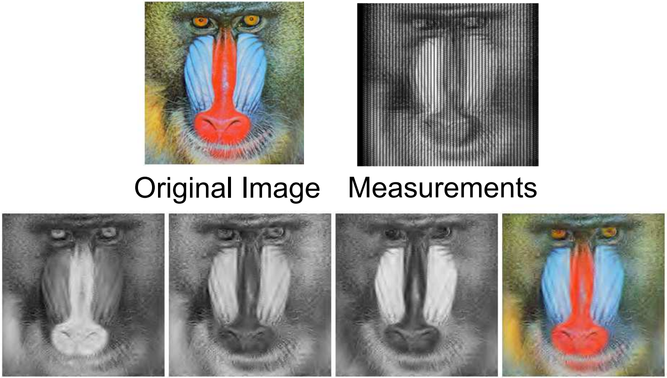

R G B Reconstructed Image

Single Shot CASSI System

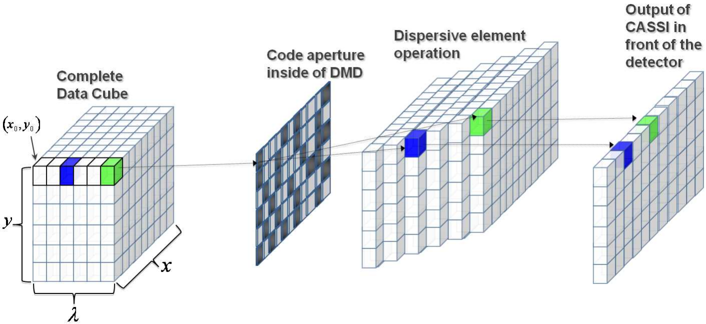

Object with spectral information only in (x_o, y_o)
Only two spectral component are present in the object

Single Shot CASSI System

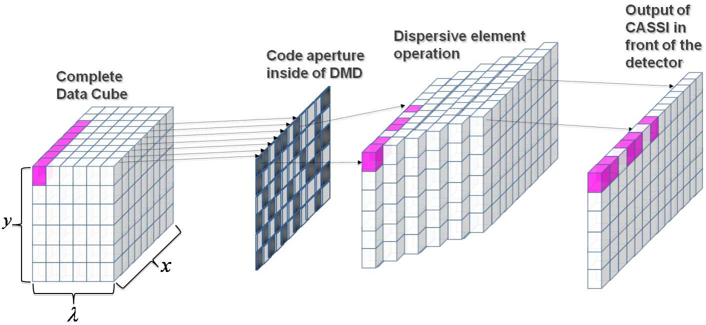

Object with spectral information only in (x_o, y_o)

Single Shot CASSI System

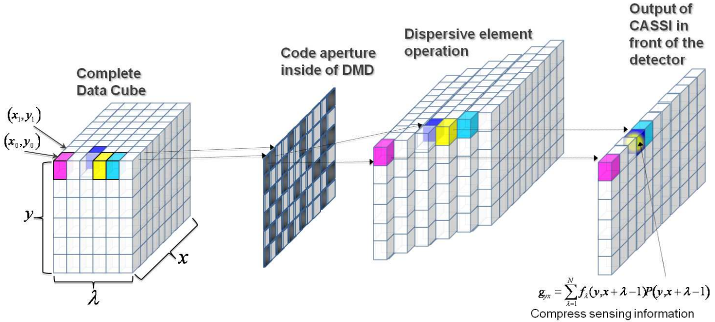

One pixel in the detector has information from different spectral bands and different spatial locations

Single Shot CASSI System

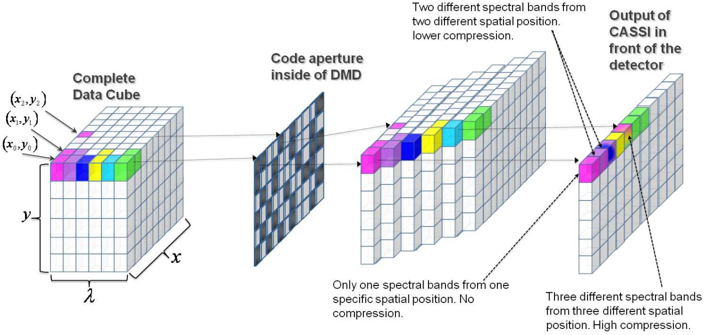

Each pixel in the detector has different amount of spectral information. The more compressed information, the more difficult it is to reconstruct the original data cube.

Single Shot CASSI System

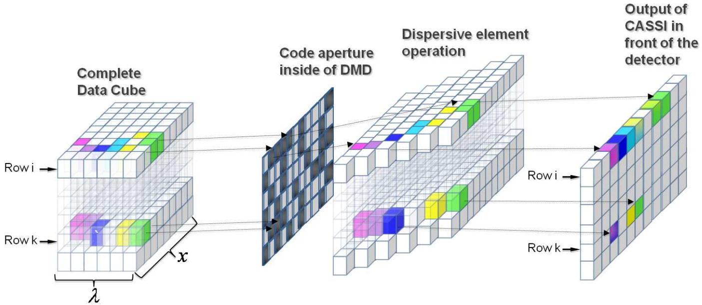

Each row in the data cube produces a compressed measurement totally independent in the detector.

Single Shot CASSI System

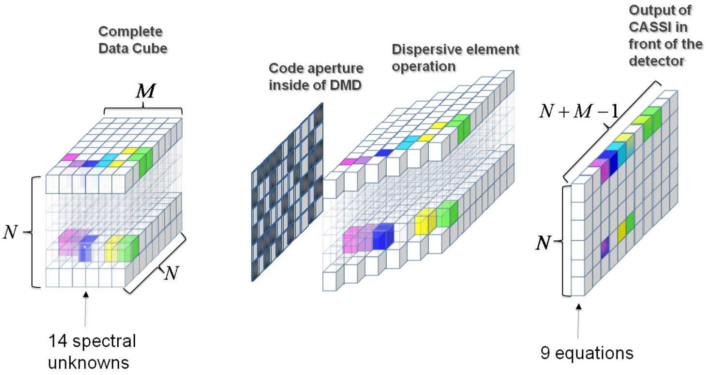

Undetermined equation system:
Unknowns $= N \times N \times M$ and Equations: $N \times (N + M - 1)$

Single Shot CASSI System

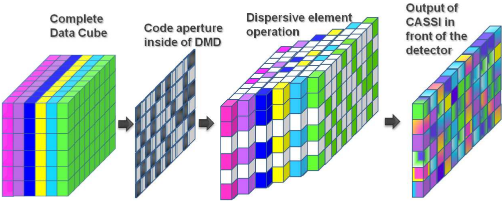

- Complete data cube 6 bands
- The dispersive element shifts each spectral band in one spatial unit
- In the detector appear the compressed and modulated spectral component of the object
- At most each pixel detector has information of six spectral components

Single Shot CASSI System

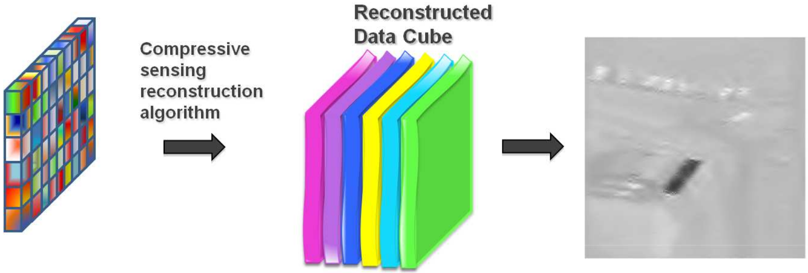

We used the $\ell_1 - \ell_s$ reconstruction algorithm †.

† S. J. Kim, K. Koh, M. Lustig, S. Boyd and D. Gorinevsky. "An interior-point method for large scale L1 regularized least squares." IEEE Journal of Selected Topics in Signal Processing, vol.1, pp. 606-617, 2007.

Coded Aperture Snapshot Spectral Image System $(CASSI)^{(a)}$

Advantages:

- Enables compressive spectral imaging
- Simple
- Low cost and complexity

Limitations:

- Excessive compression
- Does not permit a controllable SNR
- May suffer low SNR
- Does not permit to extract a specific subset of spectral bands

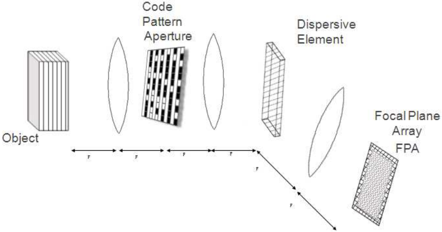

$$g_{mn} = \sum_k f_{(m+k)nk} P_{(m+k)n} + w_{nm}$$
$$= (Hf)_{nm} + w_{nm} = (HW\theta)_{nm} + w_{nm}$$

A. Wagadarikar, R. John, R. Willett, and D. Brady. "Single disperser design for coded aperture snapshot spectral imaging." Appl. Opt., Vol.47, No.10, 2008.

Bands Recovery

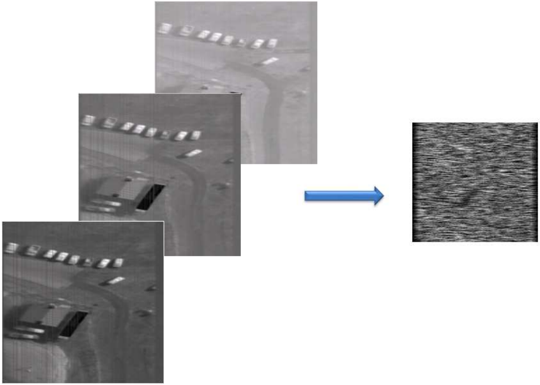

Typical example of a measurement of CASSI system. A set of bands constant spaced between them are summed to form a measurement

Multi-Shot CASSI System

- Multi-shot compressive spectral imaging system

Advantages:

- Multi-Shot CASSI allows controllable SNR
- Permits to extract a hand-picked subset of bands
- Extend Compressive Sensing spectral imaging capabilities

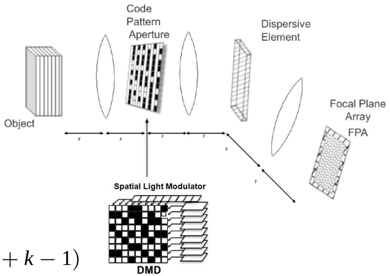

$$g_{mni} = \sum_{k=1}^{L} f_k(m, n+k-1)P_i(m, n+k-1)$$

$$= \sum_{k=1}^{L} f_k(m, n+k-1)P_r(m, n+k-1)P_g^i(m, n+k-1)$$

Ye, P. et al. "Spectral Aperture Code Design for Multi-Shot Compressive Spectral Imaging". Dig. Holography and Three-Dimensional Imaging, OSA. Apr.2010.

Mathematical Model of CASSI System

$$g_{mni} = \sum_{k=1}^{L} f_k(m, n+k-1) P_i(m, n+k-1)$$

$$= \sum_{k=1}^{L} f_k(m, n+k-1) P_r(m, n+k-1) P_g^i(m, n+k-1)$$

where i expresses i^{th} shot

Each pattern P_i is given by,

$$P_i(m, n) = P_g^i(m, n) x P_r(m, n)$$

$$P_g^i(m, n) = \begin{cases} 1 & mod(n, R) = mod(i, R) \\ 0 & \text{otherwise} \end{cases}$$

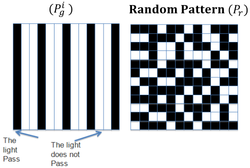

One different code aperture is used for each shot of CASSI system

Code Apertures

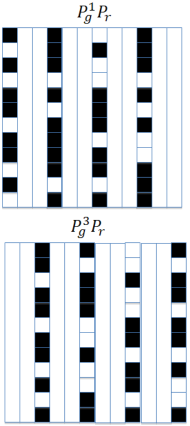

Code patterns used in multishot CASSI system

Code patterns used in multishot CASSI system

Cube Information and Subsets of Spectral Bands

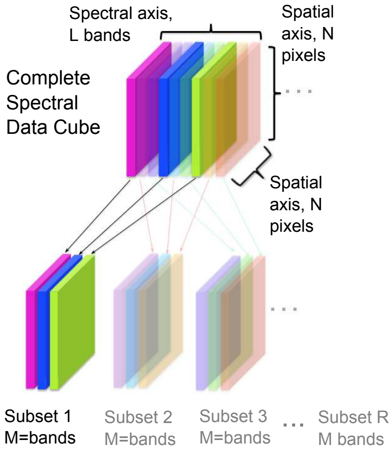

Spectral data cube $\rightarrow$ L bands R subsets of M bands each one $(L = RM)$ Each component of the subset is spaced by R bands of each other

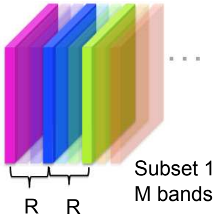

Cube Information and Subsets of Spectral Bands

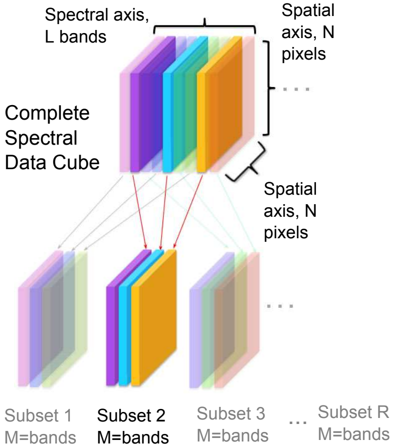

Spectral data cube $\rightarrow$ L bands R subsets of M bands each one $(L = RM)$ Each component of the subset is spaced by R bands of each other

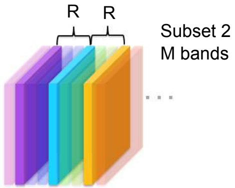

Multi-Shot CASSI System

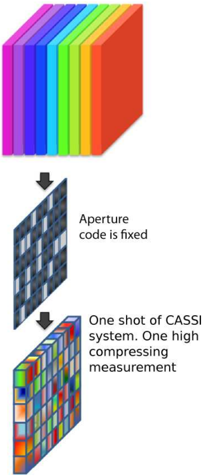

Single shot

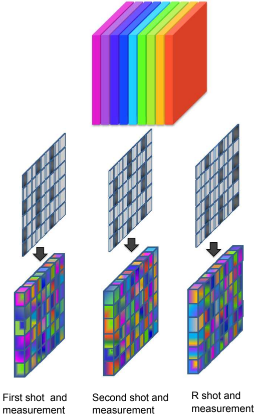

Multi-Shot

Single Shot

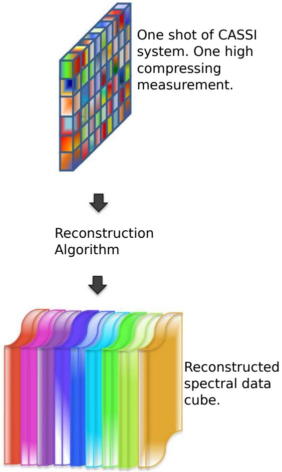

Multi-Shot

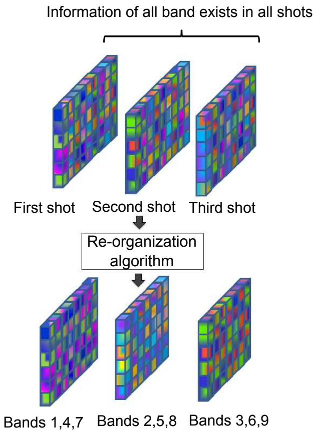

Multi-Shot

Reorder Process

$$
\begin{aligned}
g'_{mnk} &= \sum_{j=1}^{L} f_j(m, n+j-1) P_i(m, n+j-1) \\
&= \sum_{j=1}^{L} f_j(m, n+j-1) P_r(m, n+j-1) P_g^i(m, n+j-1) \\
&= \sum_{mod(n+j-1,R)=mod(i,R)} f_k(m, n+k-1) P_r(m, n+j-1) \\
&= (H_k F_k)_{mn}
\end{aligned}
$$

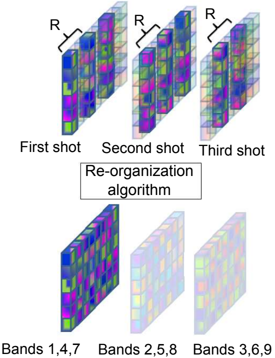

Multi-Shot

Reorder Process

$$
\begin{aligned}
g'_{mnk} &= \sum_{j=1}^{L} f_j(m, n+j-1)P_i(m, n+j-1) \\
&= \sum_{j=1}^{L} f_j(m, n+j-1)P_r(m, n+j-1)P_g^i(m, n+j-1) \\
&= \sum_{mod(n+j-1,R)=mod(i,R)} f_k(m, n+k-1)P_r(m, n+j-1) \\
&= (H_k F_k)_{mn}
\end{aligned}
$$

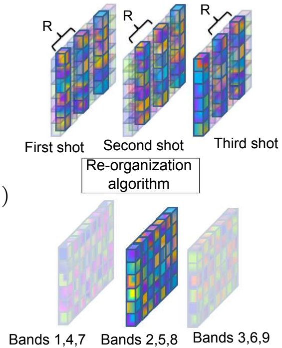

Multi-Shot

- Recover any of the subsets independently
- Recover of complete spectral data cube is not necessary

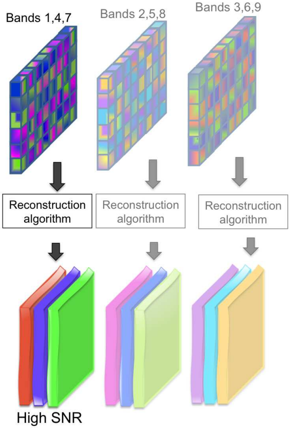

Multi-Shot

- High SNR in each reconstruction
- Enable to use parallel processing
- To use one processor for each independent reconstruction

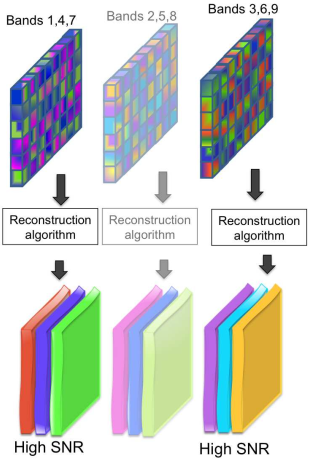

Multi-Shot

Single Shot

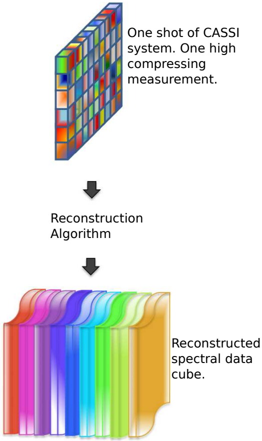

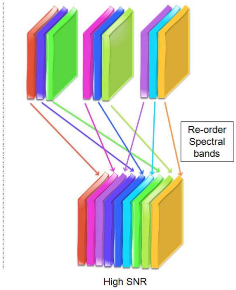

Multi-Shot Reconstruction

Reconstructed image of one spectral channel in 256*x*256*x*24 data cube from multiple shot measurements.

- (a) One shot result, $PSNR = 17.6dB$
- (b) Two shots result, $PSNR = 25.7dB$
- (c) Eight shots result, $PSNR = 29.4$
- (d) Original image

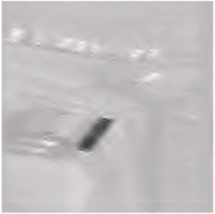

(a) One shot

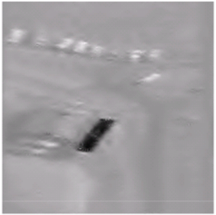

(b) 2 shots

(c) 8 shots

(d) Original

Multi-Shot Reconstruction

Reconstructed image for different spectral channels in the 256x256x24 data cube from six shot measurements.

- (a) Band 1
- (b) Band 13
- (c) Band 8
- (d) Band 20
- (a) and (b) are reconstructed from the first group of measurements
- (c) and (d) are reconstructed from the second group of measurements

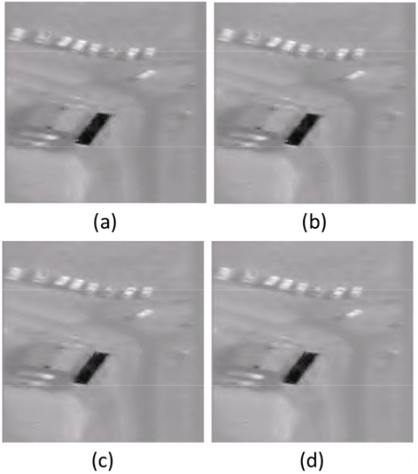

Random Convolution Imaging

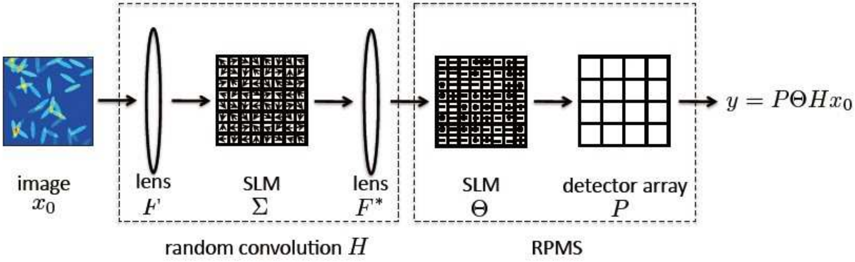

J. Romberg. "Compressive Sensing by Random Convolution." SIAM Journal on Imaging Science, July,2008.

Random Convolution Imaging

Random Convolution

Circularly convolve signal $x \in \mathbb{R}^n$ with a pulse $h \in \mathbb{R}^n$, then subsample.
The pulse is random, global, and broadband in that its energy is distributed uniformly across the discrete spectrum.

$$x * h = Hx$$

where

$$H = n^{-1/2} F^* \Sigma F$$

$$F_{t,\omega} = e^{-j2\pi(t-1)(\omega-1)/n}, 1 \leq t, \omega \leq n$$

Σ as a diagonal matrix whose non-zero entries are the Fourier transform of h.

Random Convolution

$$\Sigma = \begin{bmatrix} \sigma_1 & 0 & \cdots & \\ 0 & \sigma_2 & \cdots & \\ \vdots & & \ddots & \\ & & & \sigma_n \end{bmatrix}$$

$$\begin{aligned} \omega = 1 &: \quad \sigma_1 \sim \pm 1 \textit{ with equal probability}, \\ 2 \leq \omega < n/2+1 &: \quad \sigma_\omega = e^{j\theta_\omega}, \textit{ where } \theta_\omega \sim \textit{ Uniform}([0, 2\pi]), \\ \omega = n/2+1 &: \quad \sigma_{n/2+1} \sim \pm 1 \textit{ with equal probability}, \\ n/2+2 \leq \omega \leq n &: \quad \sigma_\omega = \sigma^*_{n-\omega+2}, \textit{ the conjugate of } \sigma_{n-\omega+2}. \end{aligned}$$

Random Convolution

Ex: if $n = 16$ $i.e.$ $x \in R^{16}$, then

$$\begin{aligned}
\sigma_1 &= 1.0000 + 0.0000i, & \sigma_2 &= -0.9998 + 0.0194i, \\
\sigma_3 &= 0.6472 - 0.7623i, & \sigma_4 &= 0.4288 + 0.9034i, \\
\sigma_5 &= -0.9211 + 0.3894i, & \sigma_6 &= 0.6110 + 0.7916i, \\
\sigma_7 &= -0.2146 + 0.9767i, & \sigma_8 &= -0.4754 + 0.8798i, \\
\sigma_9 &= 1.0000 + 0.0000i, & \sigma_{10} &= -0.4754 - 0.8798i, \\
\sigma_{11} &= -0.2146 - 0.9767i, & \sigma_{12} &= -0.6110 - 0.7916i, \\
\sigma_{13} &= -0.9211 - 0.3894i, & \sigma_{14} &= -0.4288 - 0.9034i, \\
\sigma_{15} &= -0.6472 + 0.7623, & \sigma_{16} &= -0.9998 - 0.0194i,
\end{aligned}$$

Random Convolution

H

- The action of H on a signal x can be broken down into a discrete Fourier transform, followed by a *randomization of the phase* (with constraints that keep the entries of H real), followed by an inverse discrete Fourier transform.
- Since $FF^* = F^*F = nI$ and $\Sigma\Sigma^* = I$,

$$H^*H = n^{-1}F^*\Sigma^*FF^*\Sigma F = nI$$

So convolution with h as a transformation into a random orthobasis.

Sampling at Random Locations

Simply observe entries of Hx at a small number of randomly chosen locations.
Thus the measurement matrix can be written as

$$\Phi = R_\Omega H$$

where R_Ω is the restriction operator to the set Ω (m random location subset).

Randomly Pre-Modulated Summation

- Break Hx into blocks of size n/m, and summarize each block with a single randomly modulated sum. (Assume that m evenly divides n.)
- With $B_k = \{(k-1)n/m+1, \dots, kn/m\}, k = 1, \dots, m$ denoting the index set for block k, take a measurement by multiplying the entries of Hx in B_k by a sequence of random signs and summing.

$$\phi_k = \sqrt{\frac{m}{n}} \sum_{t \in B_k} \varepsilon_t h_t$$

where h_t is the tth row of H and $\{\varepsilon_p\}_{p=1}^{n}$ are independent and take a values of ± 1 with equal probability, $\sqrt{m/n}$ is a renormalization that makes the norms of the ϕ_k similar to the norm of the h_t

Randomly Pre-Modulated Summation

The measurement matrix can be written as

$$\Phi = P\Theta H$$

where Θ is a diagonal matrix whose non-zero entries are the $\{\varepsilon_p\}$, and P sums the result over each block B_k.

Advantage

It "sees" more of the signal than random subsampling without any amplification.

Randomly Pre-Modulated Summation

$$y_{m\times 1} = \Phi_{m\times n} x_{n\times 1} = P_{m\times n}\Theta_{n\times n} H_{n\times n} x_{n\times 1}$$

where

$$P_{m\times n} = \begin{bmatrix} ones(n/m,1) & 0 & 0 & 0 \\ 0 & ones(n/m,1) & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & ones(n/m,1) \end{bmatrix}_{m\times n}$$

$$\Theta_{n\times n} = \begin{bmatrix} \pm1 & 0 & 0 & 0 \\ 0 & \pm1 & 0 & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & \pm1 \end{bmatrix}_{n\times n}$$

Randomly Pre-Modulated Summation

Why the summation must be randomly?

Imagine if we were to leave out the $\{\varepsilon_t\}$ and simply sum Hx over each B_k. This would be equivalent to putting Hx through a boxcar filter then subsampling uniformly.

Main Result

- The application of H will not change the magnitude of the Fourier transform, so signals which are concentrated in frequency will remain concentrated and signals which are spread out will stay spread out.
- The randomness of Σ will make it highly probable that a signal which is concentrated in time will not remain so after H is applied.

Main Result

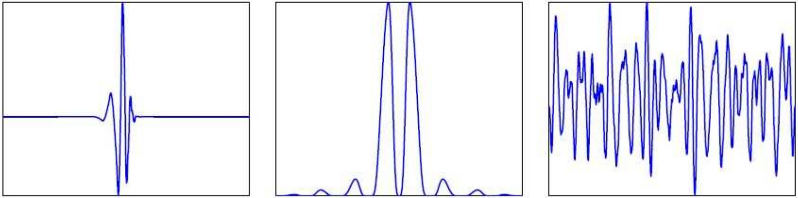

(a) (b) (c)

(a) A signal x consisting of a single Daubechies-8 wavelet.
(b) Magnitude of the Fourier transform Fx.
(c) Inverse Fourier transform after the phase has been randomized. Although the magnitude of the Fourier transform is the same as in (b), the signal is now evenly spread out in time.

J. Romberg. "Compressive Sensing by Random Convolution." SIAM Journal on Imaging Science, July,2008.

Application: Fourier Optics

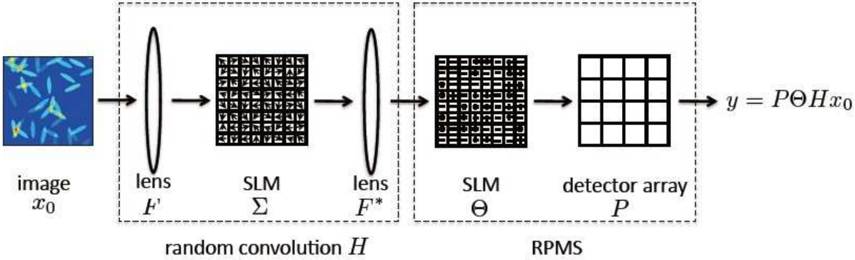

The computation $\Phi = P\Theta H$ is done entirely in analog; the lenses move the image to the Fourier domain and back, and spatial light modulators (SLMs) in the Fourier and image planes randomly change the phase.

Fourier Optics

The measurement matrix can be written as

$$\Phi = \begin{bmatrix} P \\ P\Theta H \end{bmatrix}$$

$$\min_x \mathrm{TV}(x) \quad \text{subject to} \quad \|\Phi x - y\|_2 \leq \varepsilon$$

where ε is a relaxation parameter set at a level commensurate with the noise. The result is shown in (c).

Fourier Optics

If the input signal x ($x \in R^{n\times n}$) is two dimensional like an image, *e.g.* $n = 4$, $x \in R^4$, then, in $H = n^{-1/2}F^*\Sigma F$, F is a two dimensional discrete Fourier transform instead of one dimensional, F^* is a two dimensional inverse discrete Fourier transform and

$$\Sigma = \begin{bmatrix} \sigma_{11} & \sigma_{12} & \cdots & \sigma_{1n} \\ \sigma_{21} & \sigma_{22} & \cdots & \sigma_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{n1} & \sigma_{n2} & \ldots & \sigma_{nn} \end{bmatrix}$$

where σ_ω has the conjugate relation not only in diagonal direction but also in row and column direction.

Fourier Optics

If $n = 4$, Σ can be constructed as

$$\begin{bmatrix} -1.0000 + 0.0000i & -0.4474 - 0.8944i & -1.0000 + 0.0000i & -0.4474 + 0.8944i \\ -0.2593 + 0.9658i & 0.5878 + 0.8090i & -0.1072 + 0.9942i & 0.8561 + 0.5167i \\ 1.0000 + 0.0000i & 0.9950 + 0.0995i & -1.0000 + 0.0000i & 0.9950 - 0.0995i \\ -0.2593 - 0.9658i & 0.8561 - 0.5167i & -0.1072 - 0.9942i & 0.5878 - 0.8090i \end{bmatrix}$$

Fourier Optics

In $\Phi = P\Theta H$, P sums the results over each block *e.g.* 4×4. Θ is a matrix whose entries are independent and take a values of ± 1 with equal probability.

If $n = 4$, then

$$\Theta = \begin{bmatrix} 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & 1 \\ -1 & -1 & 1 & 1 \\ 1 & -1 & -1 & -1 \end{bmatrix}$$

Fourier Optics

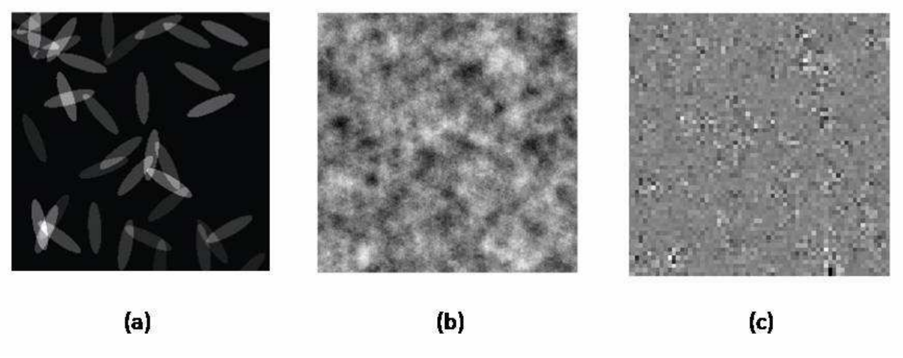

Fourier optics imaging experiment.
(a) The 256×256 image x.
(b) The 256×256 image Hx.
(c) The 64×64 image $P\theta Hx$.

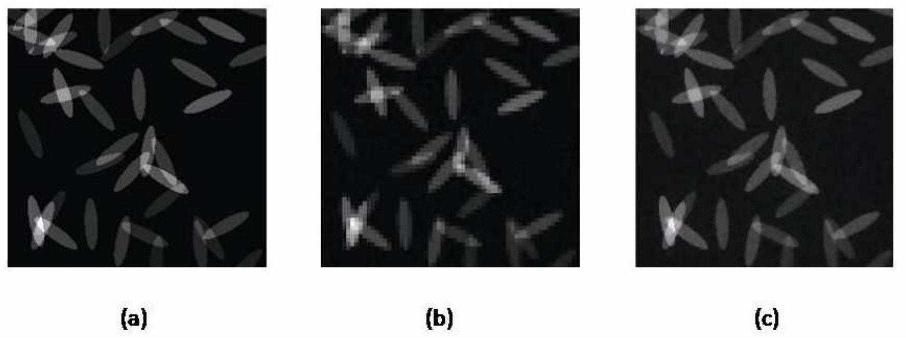

(a) The 256 $\times$ 256 image we wish to acquire.

(b) High-resolution image pixellated by averaging over 4 $\times$ 4 blocks.

(c) The image restored from the pixellated version in (b), plus a set of incoherent measurements. The incoherent measurements allow us to effectively super-resolve the image in (b).

Fourier Optics

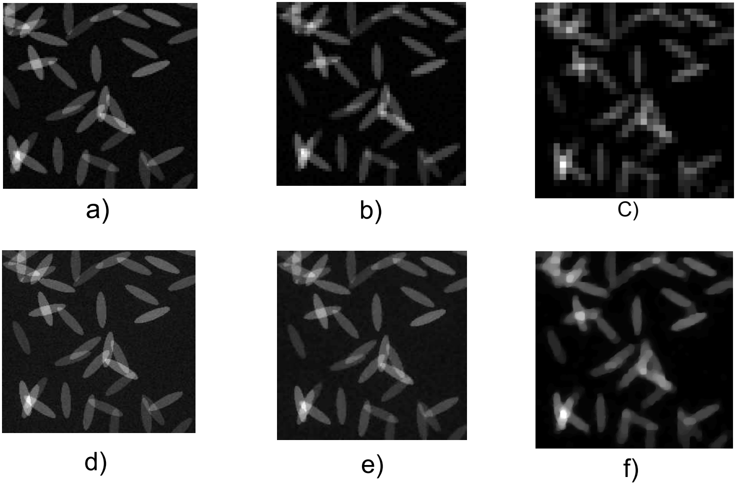

Pixellated images: (a) 2×2. (b) 4×4. (c) 8×8. Restored from: (d) 2×2 pixellated version. (e) 4×4 pixellated version. (f) 8×8 pixellated version.

Fourier Optics

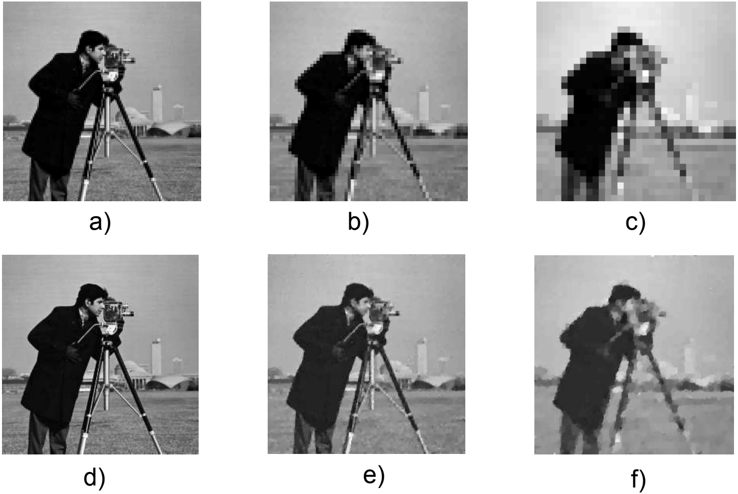

Pixellated images: (a) 2×2. (b) 4×4. (c) 8×8. Restored from: (d) 2×2 pixellated version. (e) 4×4 pixellated version. (f) 8×8 pixellated version.

Random Convolution Spectral Imaging

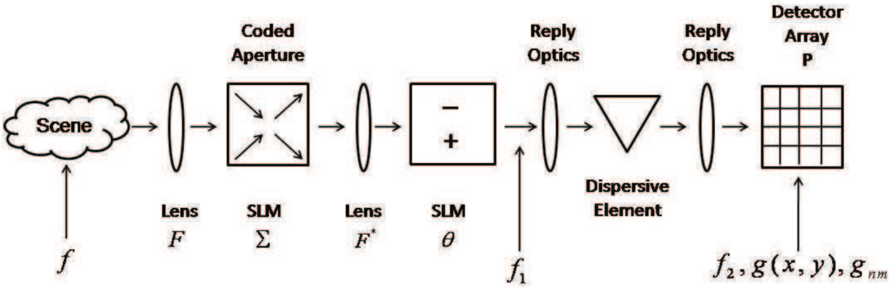

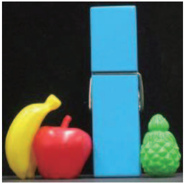

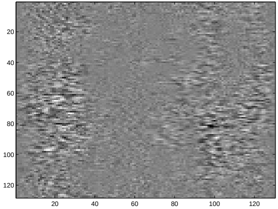
20
40
60
80
100
120
20
40
60
80
100
120

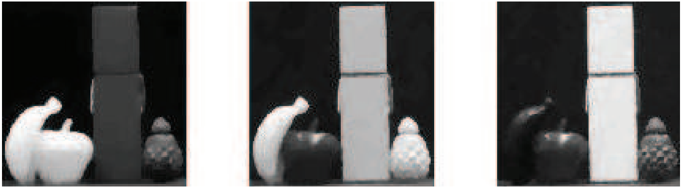

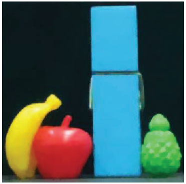

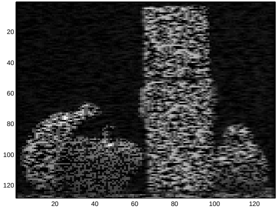
20
40
60
80
100
120
20 40 60 80 100 120