Compressive Sensing

Li-Fang Cheng

April 28th, 2014

Compressive Sensing at a Glance

- Compressive sensing (CS), also know as *compressed sensing*, *compressive sampling*, has opened a new era of signal processing … and more!

***Wired*, 2010:**

*"**Compressed sensing**, or CS, the paradigm-busting field in mathematics that's reshaping the way people work with large data sets. Only **six** years old, CS has already inspired more than **a thousand** papers and pulled in **millions of dollars** in federal grants.*

*In 2006, Candes' work on the topic was rewarded with the **$500,000 Waterman Prize**, the highest honor bestowed by the National Science Foundation. It's not hard to see why. Imagine **MRI** machines that take seconds to produce images that used to take up to an hour, **military software** that is vastly better at intercepting an adversary's communications, and **sensors** that can analyze distant interstellar radio waves. Suddenly, data becomes easier to gather, manipulate, and interpret. "*

Outline

- Theory
  - Compressive Sensing
  - Informative Sensing

- Applications
  - Image Denoising
  - Face Recognition
  - Background Subtraction & Tracking
  - Media Recovery
  - Others…

Part I:

Compressive Sensing (CS) Theory

Compressive Sensing (in one slide)

- If a signal $x_{N\times 1}$ is K-sparse in a basis / dictionary $\Psi_{N\times N}$, then it is possible to *fully recover* x by a measurement $y_{M\times 1}$ taken by $\Phi_{M\times N}$, $M \ll N$ under some conditions; i.e.

$$\text{for} \quad x = \sum_{i=1}^{N} v_i \psi_i$$

$$\text{take} \quad y = \Phi x$$

- ➢ ψ_i are basis vectors of orthonormal Ψ
- ➢ Vector v has only K nonzero elements, $K << N$
- ➢ Measurement matrix Φ satisfies *Restricted Isometry Property (RIP)*
- ➢ $M \geq cK \log(N/K) \ll N$

Compressive Sensing (step by step)

Image

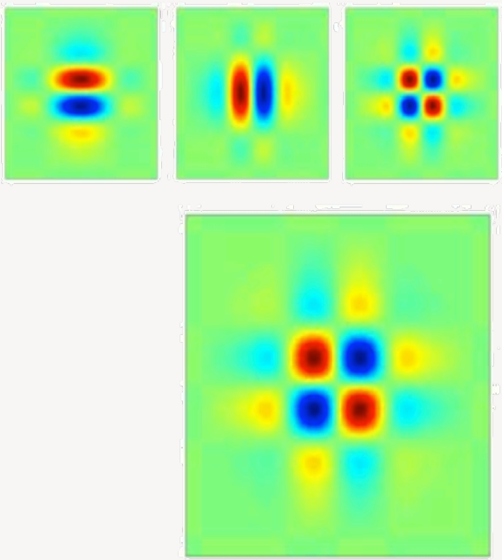

Orthonormal Basis
ex. Gabor, Wavelets, …

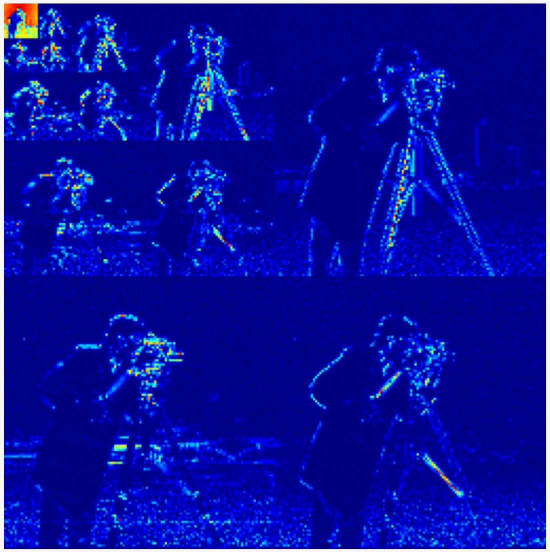

Sparse Representation

$$x = \Psi \mathbf{v}$$

K-sparse

Compressive Sensing (step by step)

- Conventional coding methods compress / reconstruct the signals through discarding sufficiently small coefficients

$$x = \Psi \mathbf{v} \approx \sum_{i \in K \text{ largest}} v_i \psi_i$$

Original (24 bit/pixel)

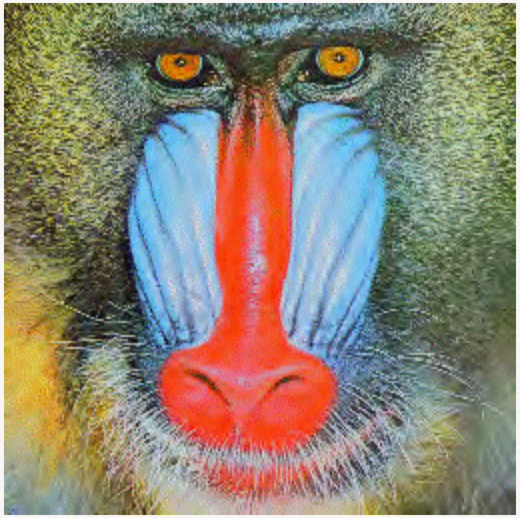

JPEG2000 (2 bit/pixel)

JPEG2000 (0.5 bit/pixel)

Compressive Sensing (step by step)

- Compressed sensing theorems state that, for signals that are sparse in some domains, they can be fully reconstruct using only few designed measurements
  - Sensing = Compressing!
  - # of required measurements depends on *sparsity*

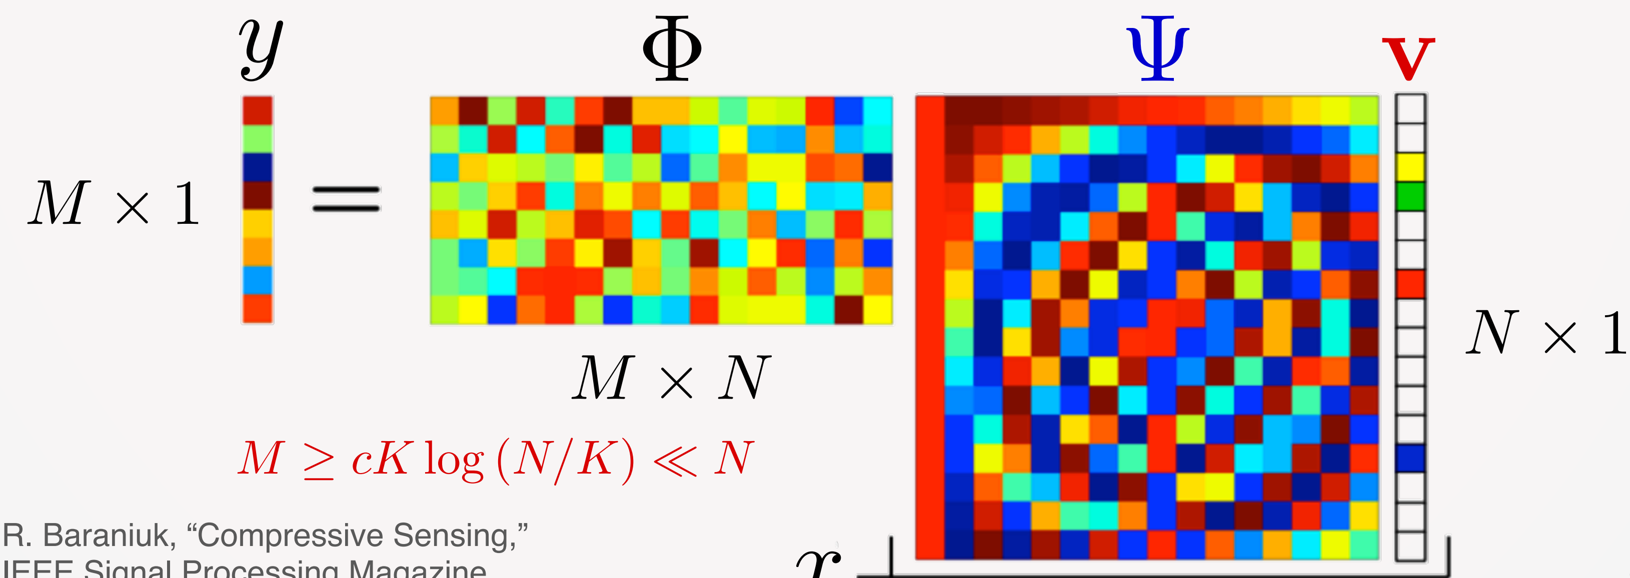

$$M \geq cK \log\left(N/K\right) \ll N$$

R. Baraniuk, “Compressive Sensing,”
IEEE Signal Processing Magazine

Compressive Sensing (step by step)

- In order to recover the sparse signal, the measurements should be taken from a measurement matrix whose product with the basis satisfies *Restricted Isometry Property (RIP)*

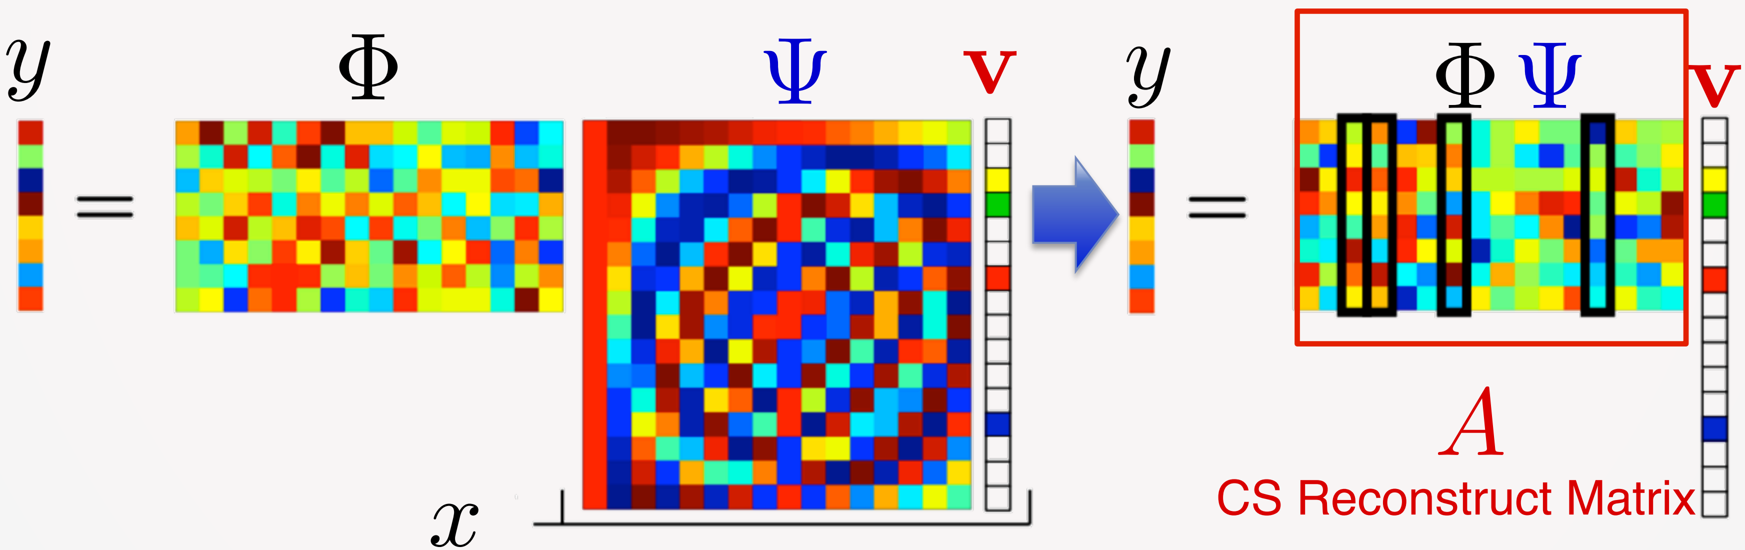

R. Baraniuk, "Compressive Sensing,"
IEEE Signal Processing Magazine

Restricted Isometry Property (RIP)

- A matrix A is said to satisfy the RIP of order K with isometry constant δ_K, which is not too close to one, such that

$$(1-\delta_K)\|v\|_2^2 \leq \|Av\|_2^2 \leq (1+\delta_K)\|v\|_2^2$$

for K-sparse vector v , or

$$(1-\delta_K)\|v_1 - v_2\|_2^2 \leq \|Av_1 - Av_2\|_2^2 \leq (1+\delta_K)\|v_1 - v_2\|_2^2$$

- In plain English!
  - A approximately preserves the Euclidean length for K-sparse signals
  - All subsets of K columns taken from A are in fact *nearly orthogonal*

E. J. Candes, "Introduction to Compressive Sampling," IEEE Signal Processing Magazine, 2008

The Measurement Matrix

- Key problem: how to design the measurement matrix Φ such that $A = \Phi\Psi$ satisfies RIP of order $2K$?
  - In practice, there is no computationally feasible way to check RIP property for a given matrix Φ

    $\because$ require $\binom{N}{K}$ checks for all K nonzero elements

- Solution: pick *random matrices*, which are usually *incoherent* with any fixed basis Ψ

- Examples
  - Random Gaussian
  - Random Bernoulli
  - Hadamard matrix

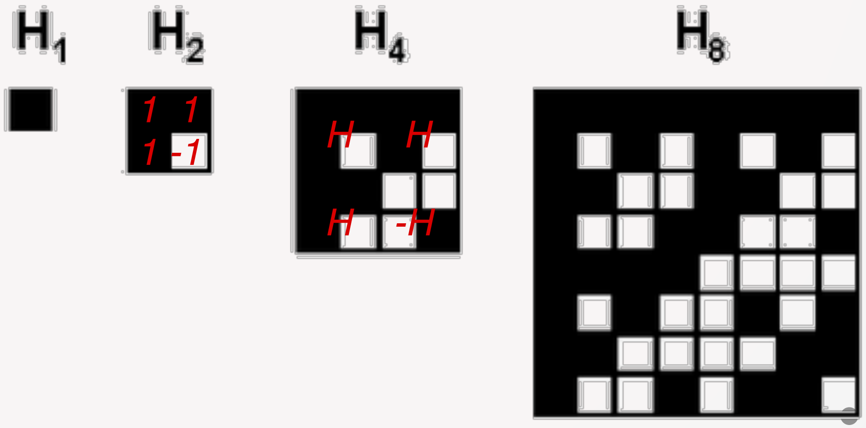

CS Recovery

- Reconstruction / Decoding Problem

$$\overbrace{y = \Phi}^{\text{Given: measurements, and sensing matrix}} \underbrace{\Psi}_{\text{Assumption: Sparsifying Basis}} \underbrace{\mathbf{v}}_{\text{Target: Sparse Representation}}$$

  - $\because \dim(y) \ll \dim(x)$, this is an ill-posed problem in general

CS Recovery Algorithms

- Signal reconstruction algorithm aims to find signal's sparse coefficient vector
    - L2 norm (energy) minimization:

$$\min \|v\|_2 \text{ subject to } Av = y$$

$$v = (A^T A)^{-1} A^T y$$

$\Rightarrow$ pseudo inverse is closed form, but does not find sparse solution

    - L0 norm minimization:

$$\min \|v\|_0 \text{ subject to } Av = y$$

$\Rightarrow$ can recover K-sparse signal exactly with high probability using only $M = K + 1$ measurements, but solving it is *NP-complete* (need exhaustive search all $\binom{N}{K}$ positions nonzero coefficients)

R. Baraniuk, "Compressive Sensing," IEEE Signal Processing Magazine

CS Recovery Algorithms

- Signal reconstruction algorithm aims to find signal's sparse coefficient vector
  - L1 norm minimization:

$$\min \|v\|_1 \text{ subject to } Av = y$$

$\Rightarrow$ can exactly recover K-sparse signals with high probability using only $M \geq cK \log(N/K) \ll N$ measurements (*mild oversampling*)

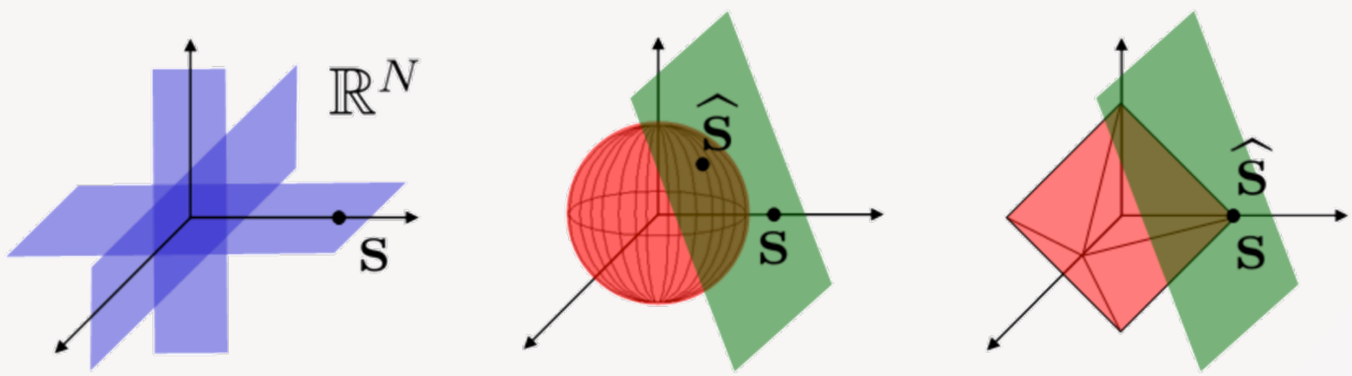

R. Baraniuk, "Compressive Sensing," IEEE Signal Processing Magazine

CS for Image Reconstruction

An one megapixel image, with only *25K* nonzero wavelet coefficients

CS for Image Reconstruction

An one megapixel image, with only *25K* nonzero wavelet coefficients

L1-reconstruction using *96K* random measurements

L1 Recovery Algorithms

- L1-mininization, a.k.a Basis Pursuit (BP)

$\boxed{\text{Min-L1 with equality constraints}}$

$$\min \|v\|_1 \text{ subject to } Av = y$$

$\boxed{\text{Min-L1 with quadratic constraints}}$

$$\min \|v\|_1 \text{ subject to } \|Av - y\|_2 \leq \epsilon$$

  - Solvable in polynomial time $O(N^3)$ by linear programming (LP)
  - Methods supported by L1-magic: Primal-dual algorithm (LP), log-barrier algorithm (SOCP), Newton iterations

http://users.ece.gatech.edu/~justin/l1magic/
S. S. Chen and D. L. Donoho, 1994

L1 Recovery Algorithms

- If the underlying signal is a 2D image, an alternative recovery model is that the *gradient* is sparse, i.e. solving by *total variation (TV) minimization*

$$\begin{aligned}\mathrm{TV}(v) &= \textstyle\sum_{i,j} \sqrt{(v_{i+1,j} - v_{i,j})^2 + (v_{i,j+1} - v_{i,j})^2} \\ &= \textstyle\sum_{i,j} \|D_{i,j} v\|_2\end{aligned}$$

$$\min \ \mathrm{TV}(v) \text{ subject to } \|Av - y\|_2 \leq \epsilon$$

$$\boxed{\begin{array}{ll} & \min_{u,v} \sum_{i,j} u_{i,j} \\ \text{subject to} & \|D_{i,j} v\|_2 \leq u_{i,j} \\ & \|Av - y\|_2 \leq \epsilon, \quad \forall i,j \end{array}}$$

Second-Order Cone Problem (SOCP)

http://users.ece.gatech.edu/~justin/l1magic/

Boyd et. al, 1997

Phantom Recovery using L1-magic

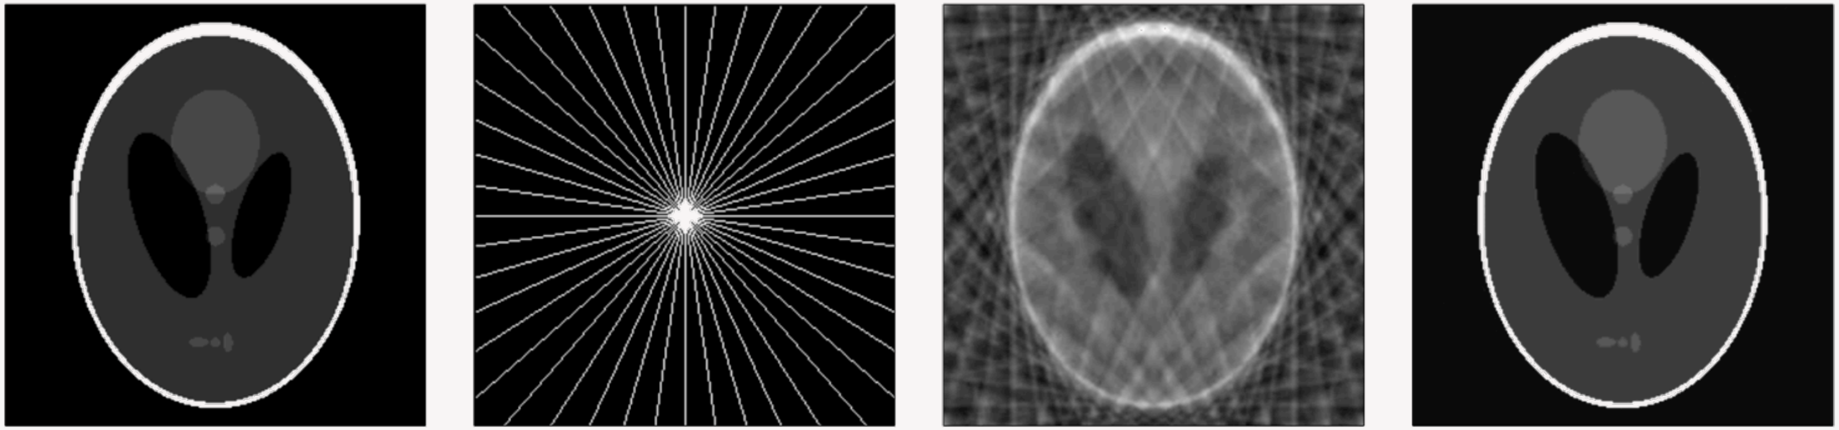

(a) Phantom (b) Sampling pattern (c) Min energy (d) min-TV reconstruction

http://users.ece.gatech.edu/~justin/l1magic/

L1 Recovery Algorithms

- Matching Pursuit (MP): OMP, StOMP, CoSaMP,…
    - Iterative greedy algorithms for basis
    - Faster, usually in time $O(KMN)$, and easier to implement
    - Many variations

$f(\mathbf{x})=\|\mathbf{y}-\mathbf{A}\mathbf{x}\|_2^2 \Rightarrow$ CoSaMP (Compressive Sampling MP) [Needell and Tropp]

State Variables: Signal estimate, $\hat{\mathbf{x}}$ support estimate: T

Initialize estimate, residual and support

$$\hat{\mathbf{x}}=0,\ T=\mathrm{supp}(\hat{\mathbf{x}}),\ \mathbf{r}=\mathbf{y}$$

Correlate residual with dictionary → signal proxy

$\mathbf{A}^T$ $\mathbf{r}$ $=$ $\mathbf{p}$

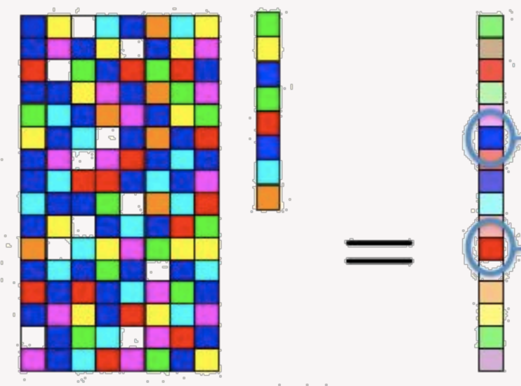

$$r = y - A\widehat{x}$$

$$\langle a_k, r\rangle = p_k$$

Select location of largest *2K* correlations

$\mathrm{supp}(\mathbf{p}|_{2K})$

Add to support set

$$\Omega = \mathrm{supp}(p|_{2K}) \cup T$$

Invert over support

$$b = A_\Omega^\dagger y$$

Truncate and compute estimate

$$T = \mathrm{supp}(b|_K)$$

$$\widehat{x} = b|_K$$

Iterate using current estimate

http://www.boufounos.com/research/sparse-function-minimization/

Comparison of Recovery Algorithms

Original Image
(512 x 512)

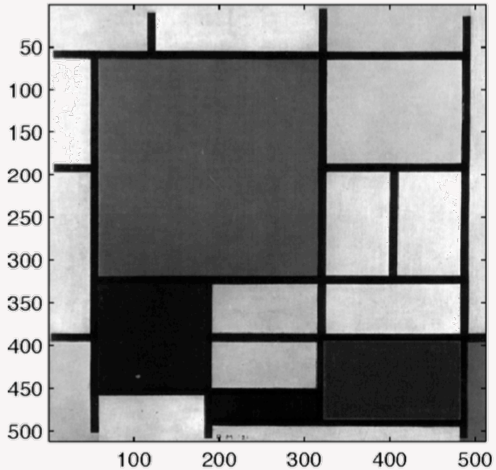

Basis Pursuit
(30 hours)

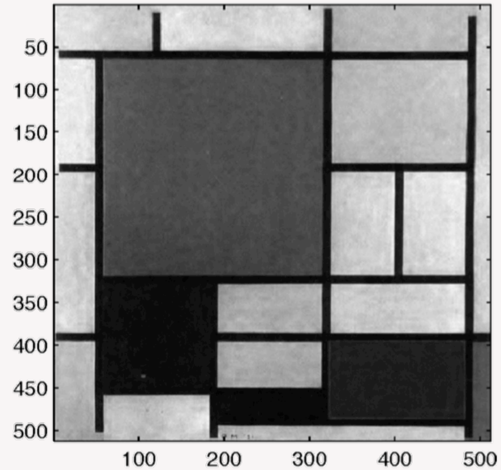

StOMP
(lower quality,
16 s)

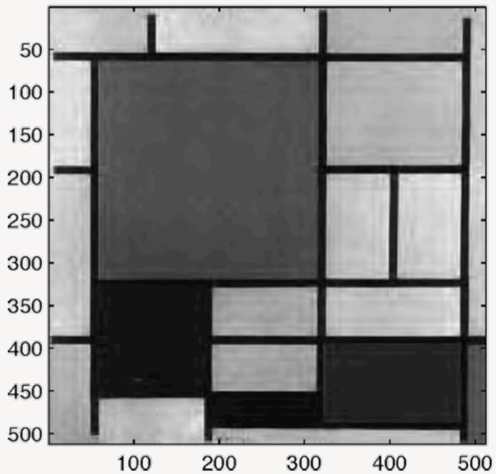

StOMP
(higher quality,
64 s)

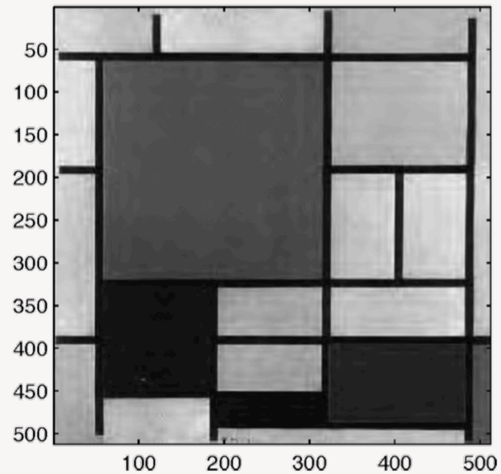

D. L. Donoho, “Sparse solution of underdetermined linear equations
by stagewise orthogonal matching pursuit,” 2006

Idea for Cost-Effective Camera

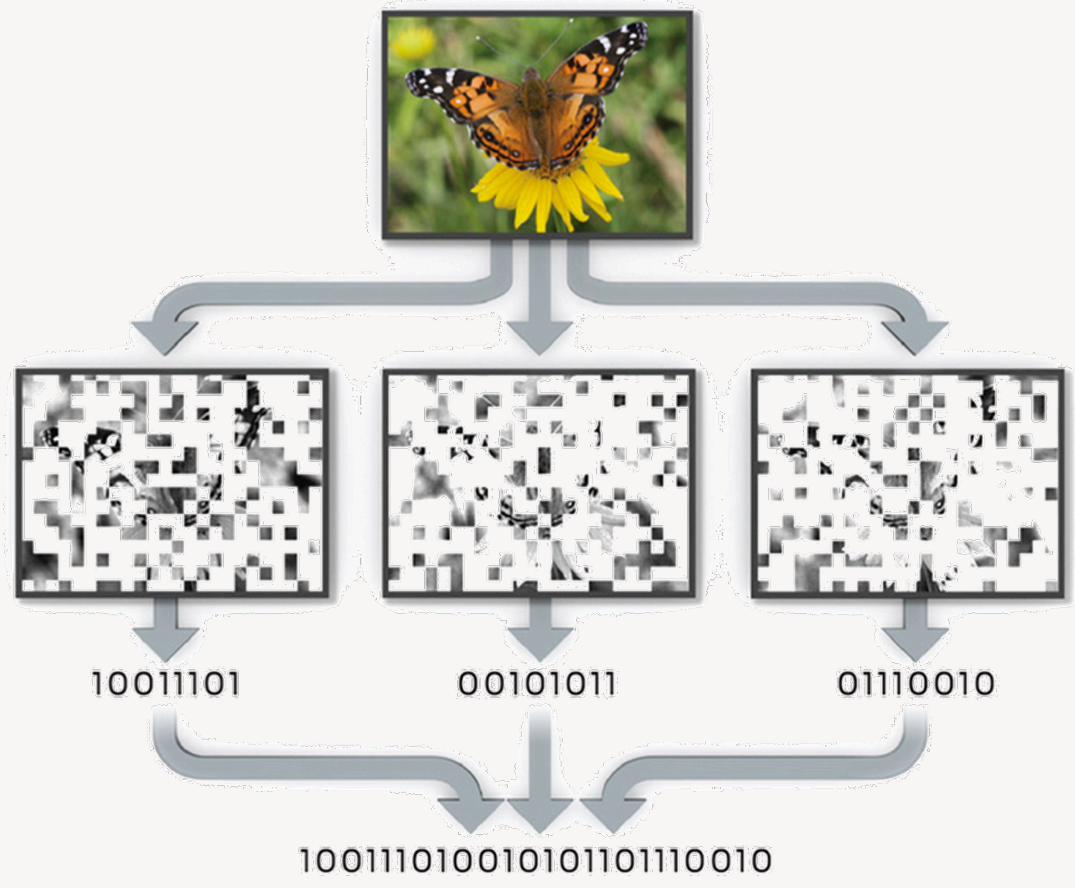

CS Single-Pixel Camera (SPC)

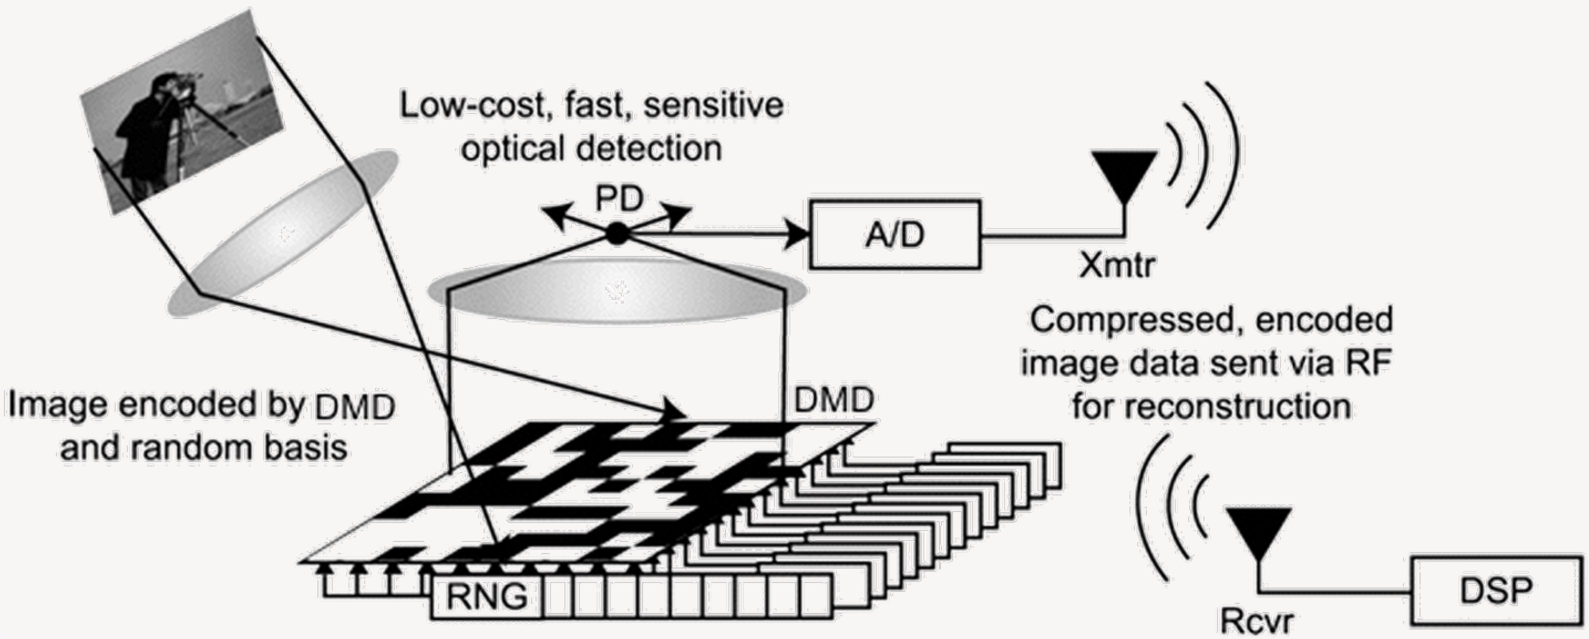

- This new camera can obtain an image or video with a *single detection element* while measuring the scene fewer times than the number of pixels / voxels

Wakin and et al, ICIP 2006
http://dsp.rice.edu/cscamera

SPC Setting

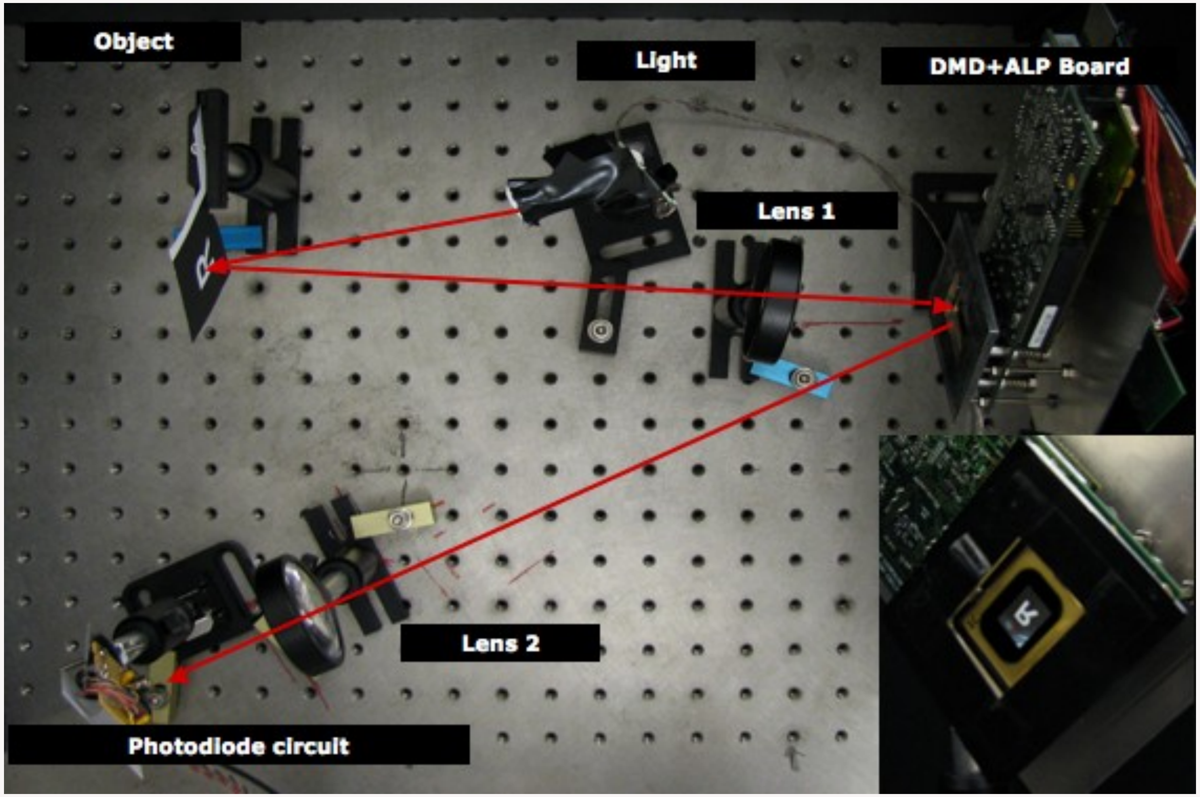

- http://dsp.rice.edu/cscamera

Architecture of DMD

- Digital Micromirror Device (DMD) by Texas Instruments

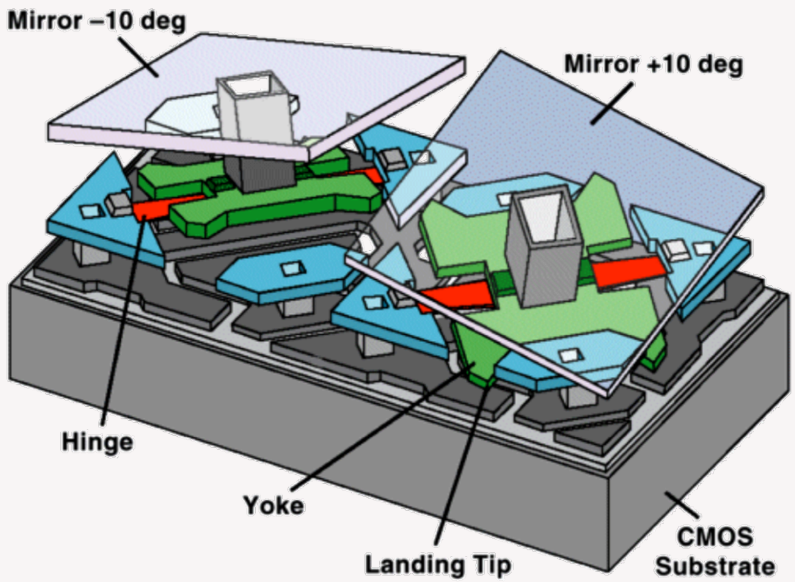

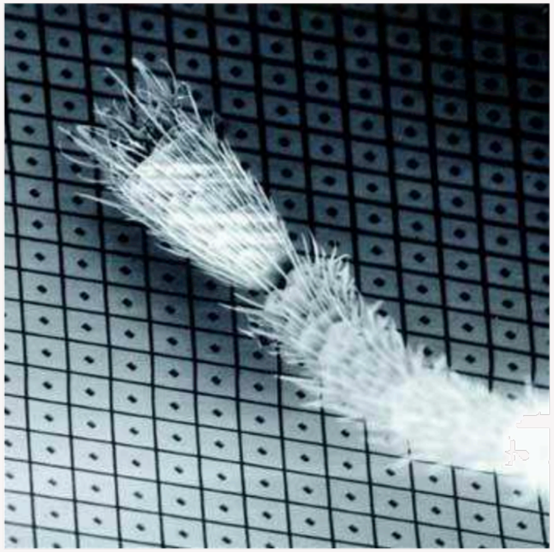

Image Acquisition Results

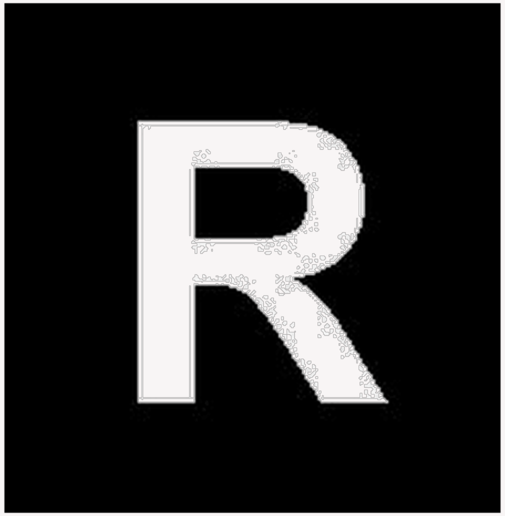

Original
(16384 pixels)

1600 measurements
(10%)

3300 measurements
(20%)

- http://dsp.rice.edu/cscamera

Image Acquisition Results

Original
(4096 pixels)

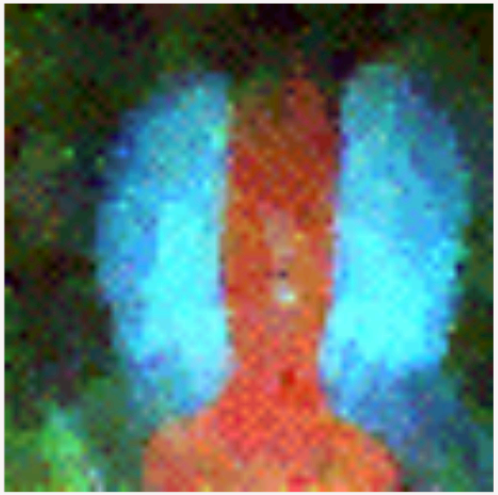

800 measurements
(20%)

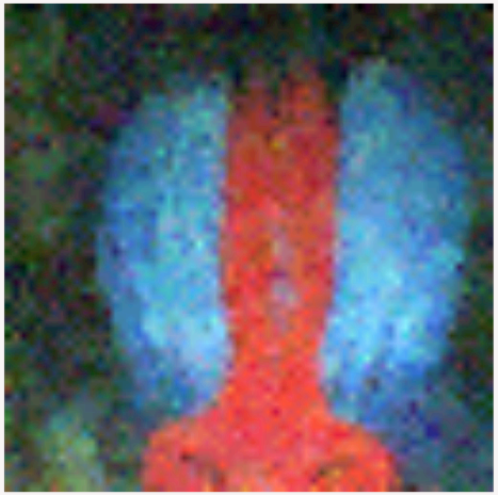

1600 measurements
(40%)

Performance

- Comparisons with other scanning methodologies

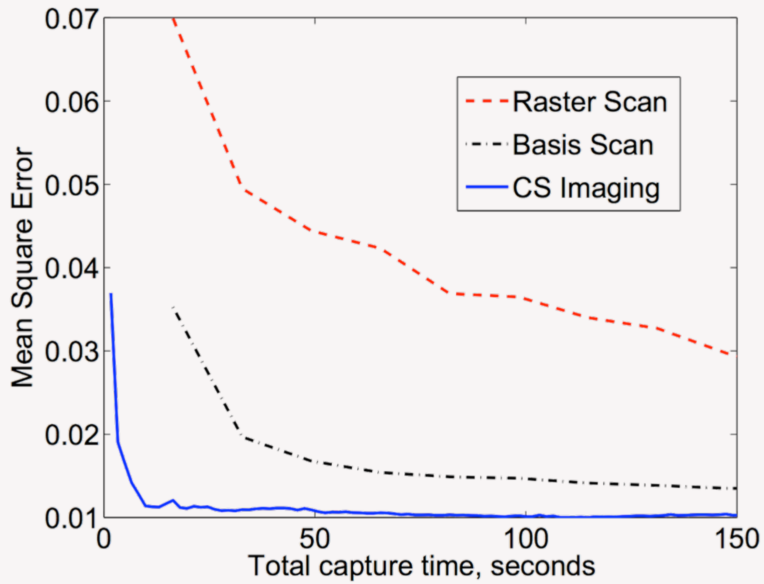

	Pixel Array	Raster Scan	Basis Scan	Compressive Sampling
Number of measurements	N	N	N	$M \leq N$
Dynamic range	D	D	$\frac{ND}{2}$	$\frac{ND}{2}$
Quantization (total bits)	NB	NB	$N(B + \log_2 N)$	$M(B + \log_2 N + \log_2 C_N + 1)$
Photon counting MSE	$\frac{P}{T}$	$N\frac{P}{T}$	$(3N-2)\frac{P}{T}$	$< 3C_N^2 M \frac{P}{T}$

- http://dsp.rice.edu/cscamera

What's Next

- CS theorem provides guidelines for perfect reconstruction, which seems not to be the case here. *Why?*

10% random measurements, *perfect* reconstruction

10% random measurements, *low quality* reconstruction

Part II:

Informative Sensing

Hyun Sung Chang, Yair Weiss, William T. Freeman, 2009

Example for Perfect Reconstruction

- Descriptions from L1-magic

- ➢ This 1 megapixel image has a *perfectly sparse* wavelet expansion
- ➢ It is a *superposition* of 25,000 Daubechies-8 wavelets
- ➢ In real world, signals and images are *not* perfectly sparse, and there is always some degree of uncertainty in the measurements

CS and Natural Images

Original Image

CS: 15% random measurements

PSNR = 25.25 dB

Slide modified from Yair Weiss's talk

CS and Natural Images

Original Image

CS: 25% random measurements

PSNR = 28.02 dB

Slide modified from Yair Weiss's talk

CS and Natural Images

Original Image

11% subsampling +
linear interpolation

PSNR = 30.03 dB

CS and Natural Images

Original Image

25% subsampling +
linear interpolation

PSNR = 33.52 dB

CS and Natural Images

- CS theory holds for ideal sparse signals, *not* natural images

Subsampling + Interp

Compressive Sensing

Weiss et al. 2006, Romberg 2008, Lustig et al. 2007,
Haupt and Nowak 2006, Seeger and Nikisch 2008

Slide modified from Yair Weiss's talk

CS and Natural Images

- CS theory holds for ideal sparse signals, *not* natural images

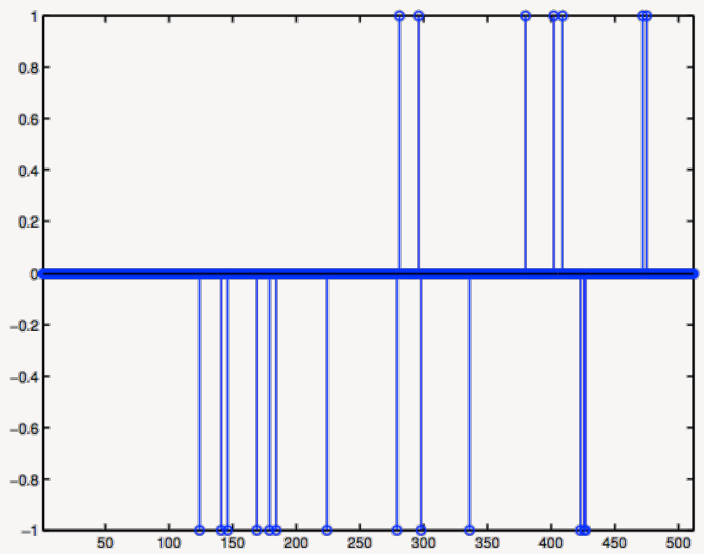

Ideal Sparse Signals

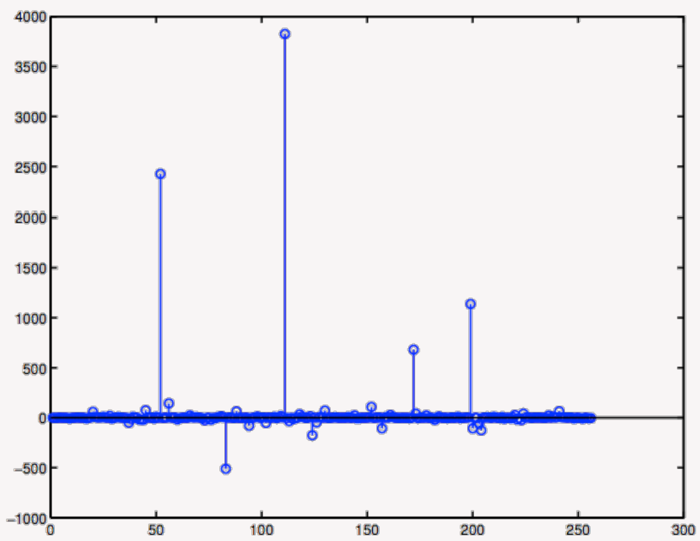

Wavelets of natural image

Natural Image Statistics

Ideal Sparse Signal | Image Wavelets | Image Wavelets (sorted)

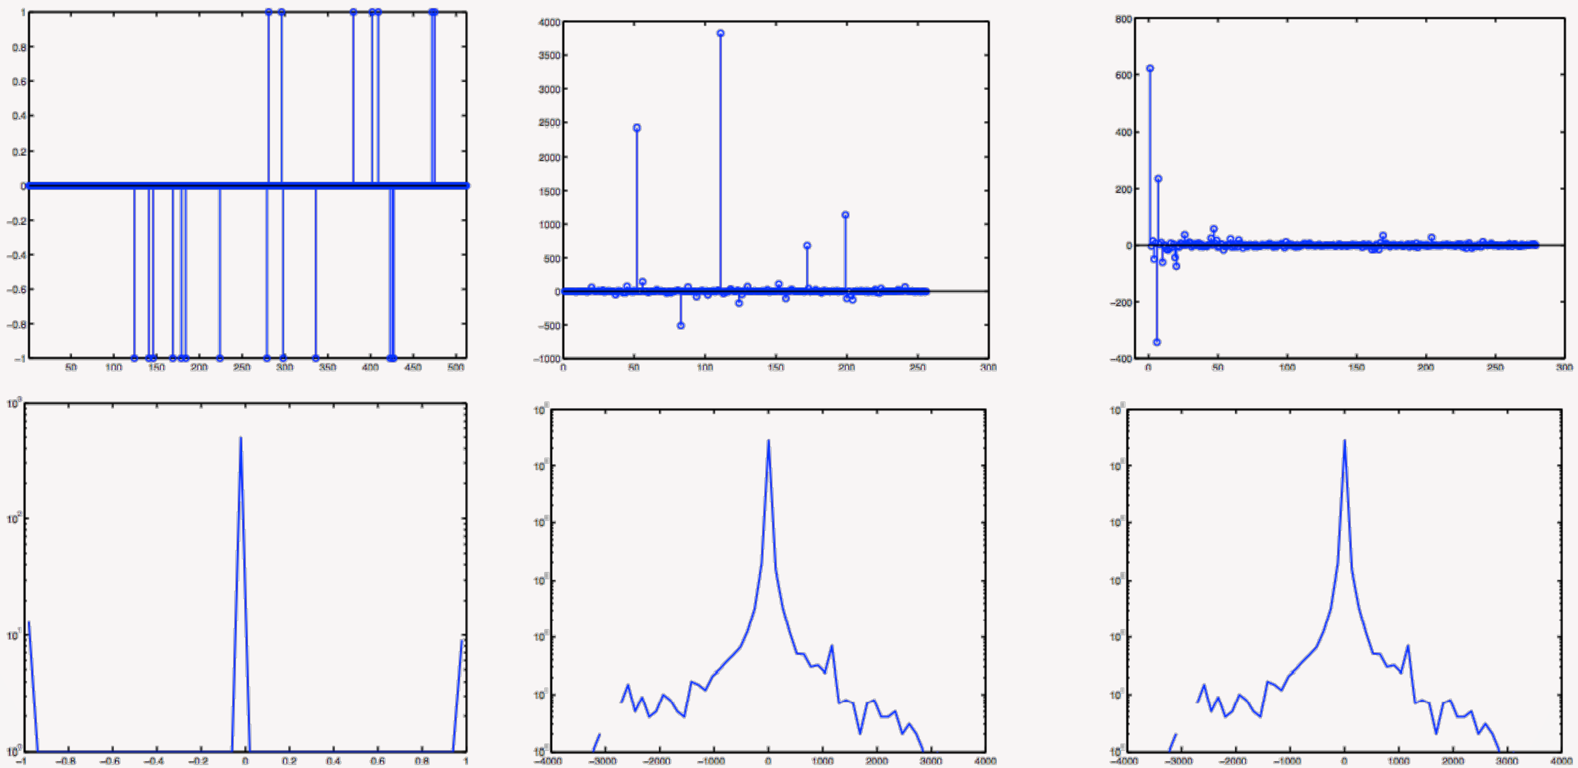

- Wavelet distribution is sparse (*statistically*)
- Decreasing variance when sorted by spatial frequency

Two Conceptual Examples

Gaussian with decreasing variance

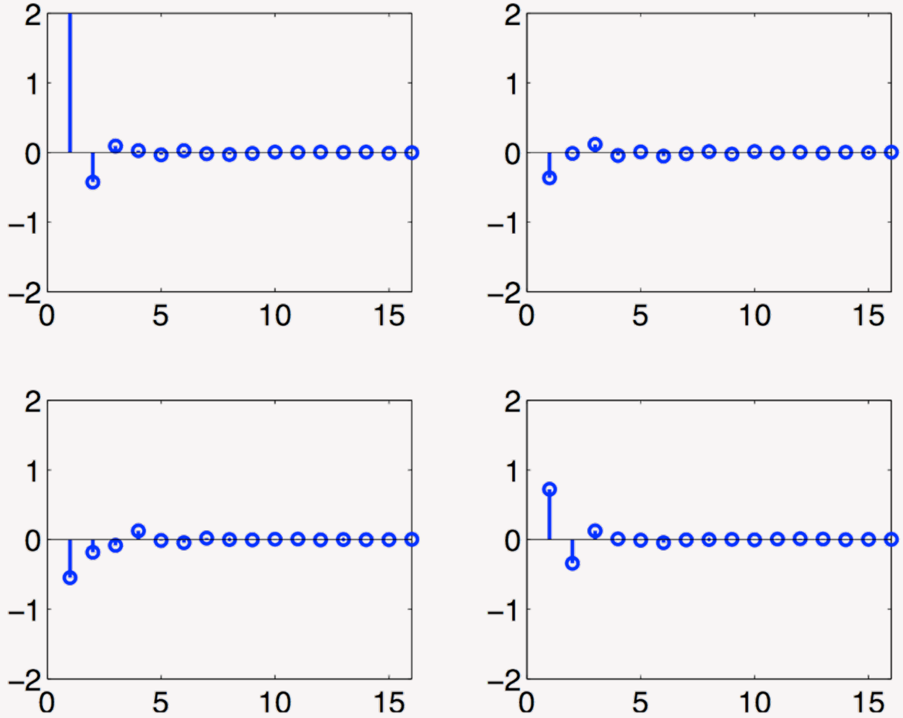

Principal Component Analysis (PCA)
gives best projections

Two Conceptual Examples

Sparse with decreasing variance

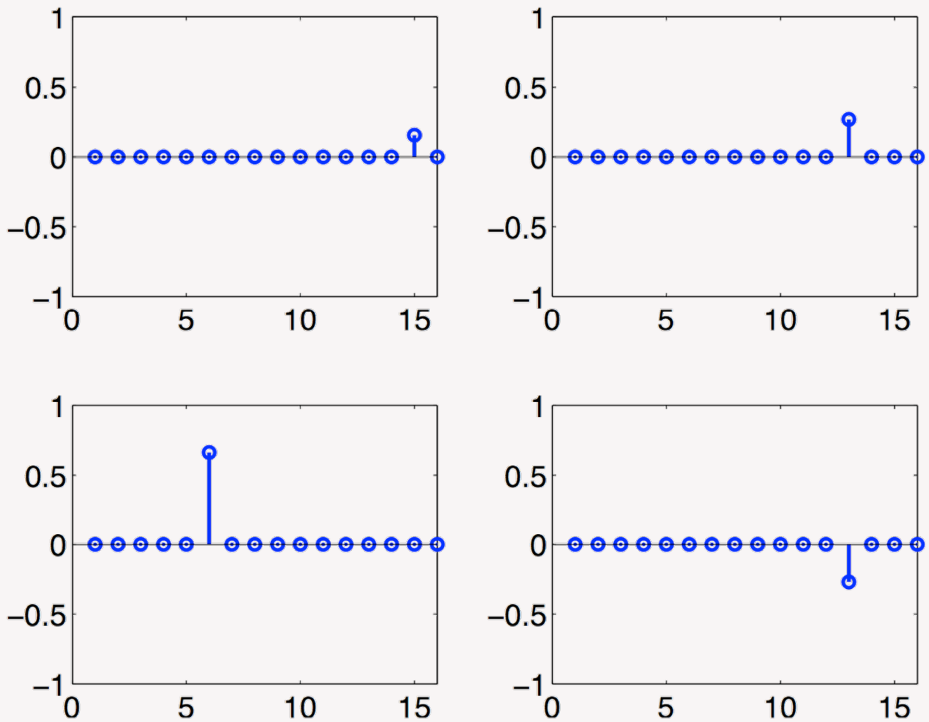

Random projections are the best
(give perfect reconstruction!)

General Case

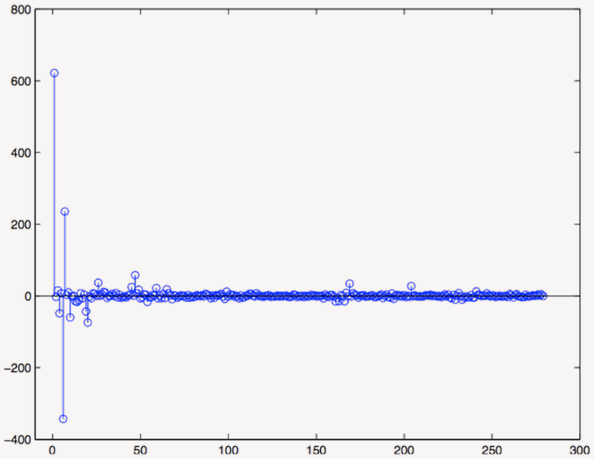

- For general distributions, *neither* PCA nor random projections are optimal
- The higher the sparsity, the higher the non-Gaussianity
- We can use entropy (negentropy J_X in this work) to characterize this property and find a better projection

Slide modified from Yair Weiss's talk

Informative Sensing

- Optimize *sensing matrix* to get maximum information
  - Given: known *distribution* over signals and a *limited* number of measurements
  - Find: best *measurements* for that distribution
  - Method: reduce the "uncertainty", or conditional entropy of unmeasured terms

$$y = Wx$$

$$\begin{bmatrix} \mathbf{y} \\ \mathbf{z} \end{bmatrix} = \begin{bmatrix} \mathbf{W}\mathbf{x} \\ \mathbf{W}_\perp \mathbf{x} \end{bmatrix} = \begin{bmatrix} \mathbf{W} \\ \mathbf{W}_\perp \end{bmatrix} \mathbf{x} = \mathbf{U}\mathbf{x}$$

$$h(\mathbf{U}\mathbf{x}) = h(\mathbf{y}, \mathbf{z}) = h(\mathbf{y}) + h(\mathbf{z}|\mathbf{y})$$

Conditional Entropy

Informative Sensing

- Optimize *sensing matrix* to get maximum information
  - Given: known *distribution* over signals and a *limited* number of measurements
  - Find: best *measurements* for that distribution
  - Method: reduce the "uncertainty", or conditional entropy of unmeasured terms

$$y = Wx$$

$$W^* = \arg \max_{W : tr(W^T W) = p} h(Ux)$$

$$h(\mathbf{y}) = h(\tilde{\mathbf{y}}) + \frac{1}{2} \ln \det(\mathbf{\Sigma}_{\mathbf{y}})$$

Bandwise (BW) Random Projection

- Use the profiling of *net capacity* to determine how many samples to take from each band
  - This distribution is related to the *sparsity* (non-Gaussianity) of the image

$$\nu(k) \stackrel{\text{def}}{=} E[h(y_1, \ldots, y_k) - h(y_1, \ldots, y_{k-1})]$$

$$\approx c_2 - \frac{2(k-1)}{n-1}(c_2 - c_\alpha) + \ln \sigma$$

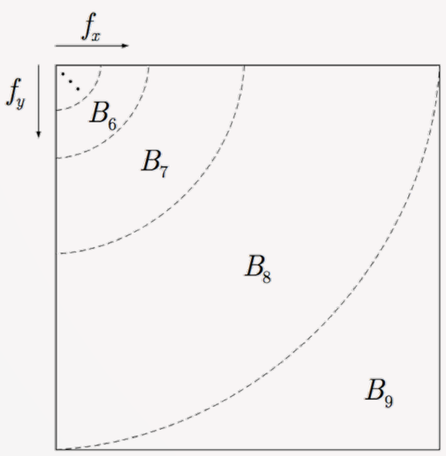

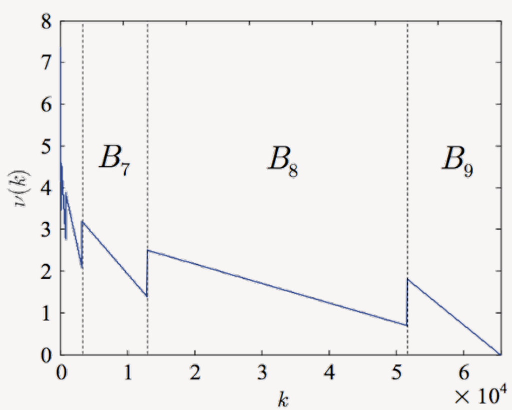

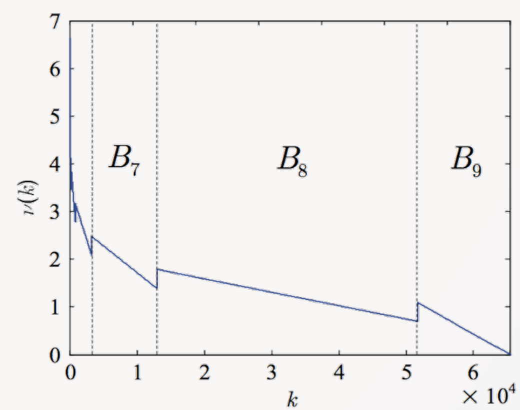

Bandwise (BW) Random Projection

- (1) Estimate and sort the net capacity in a decreasing order
- (2) Take random projections sequentially according to the total budget of number of projections

$$\mathbf{W} = \begin{bmatrix} \mathbf{W}_0 & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_1 & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{W}_{\ell_{\max}} \end{bmatrix} \mathbf{V}^{\mathsf{T}}$$

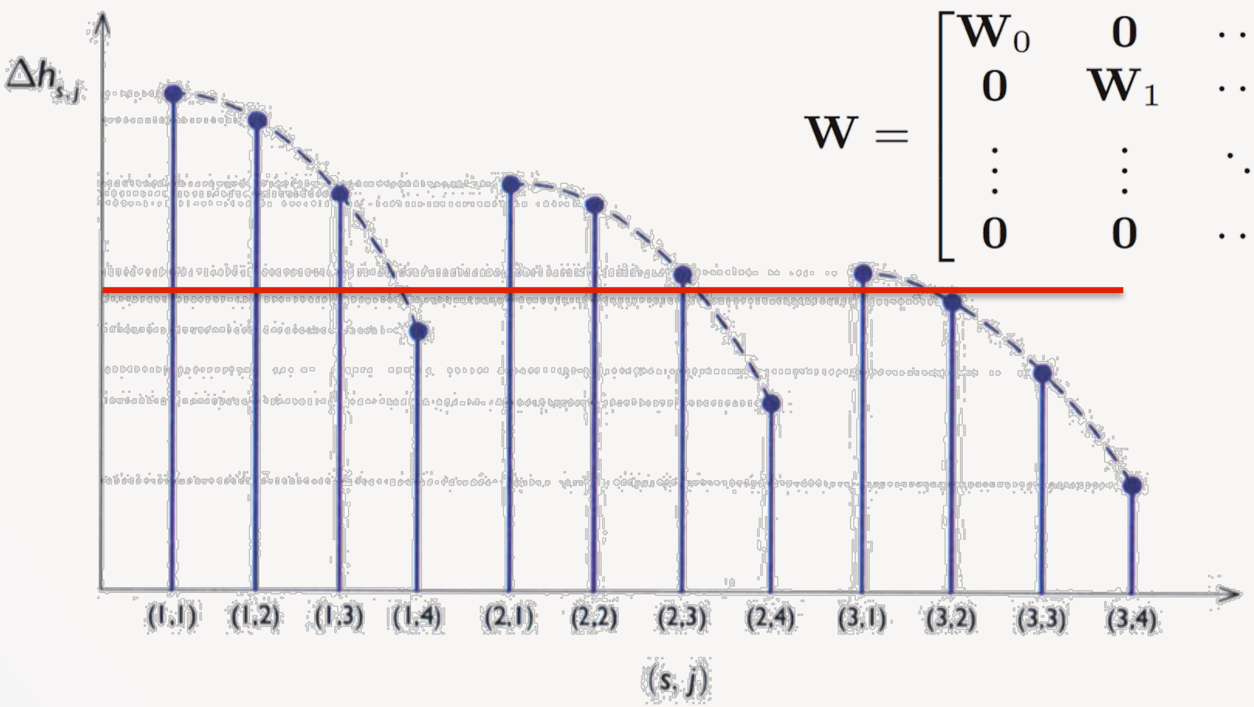

Experimental Results

- Comparison of Two Examples

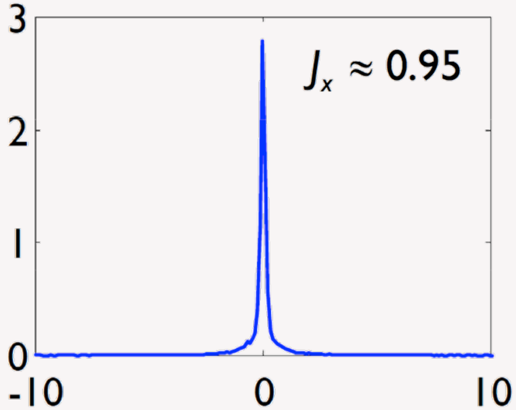

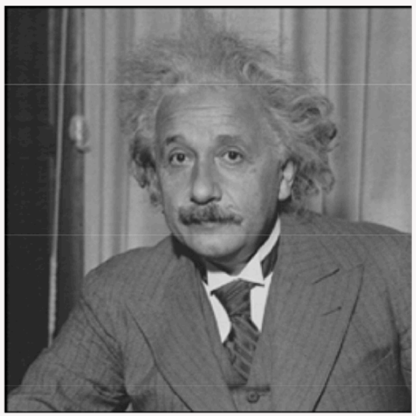

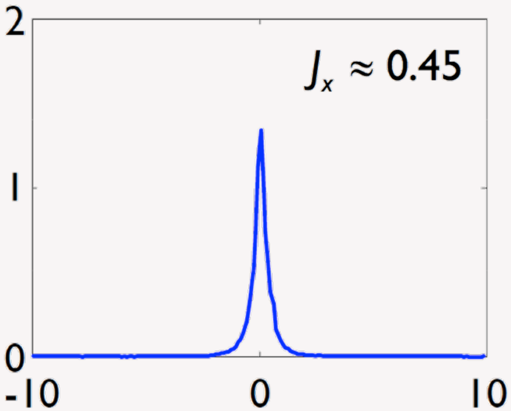

Slide modified from Yair Weiss's talk

Experimental Results

- Comparison of Two Examples

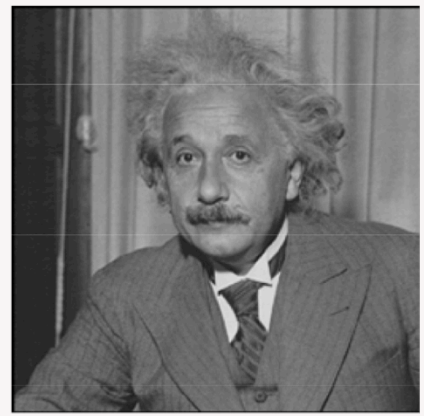

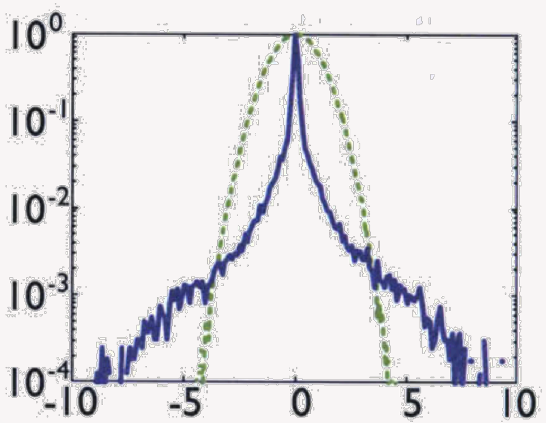

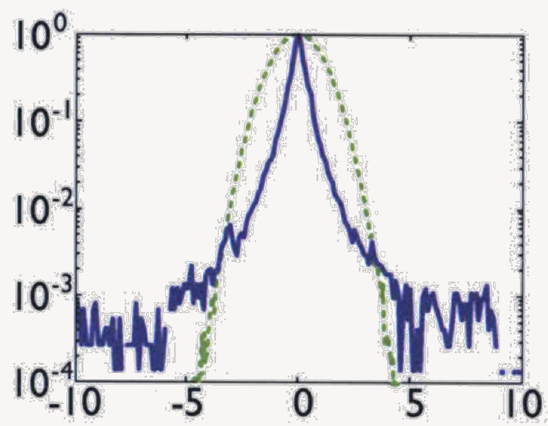

Density of Wavelet Coefficients

Experimental Results

- Comparison of Two Examples

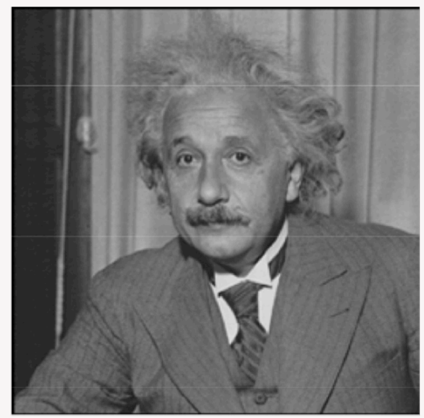

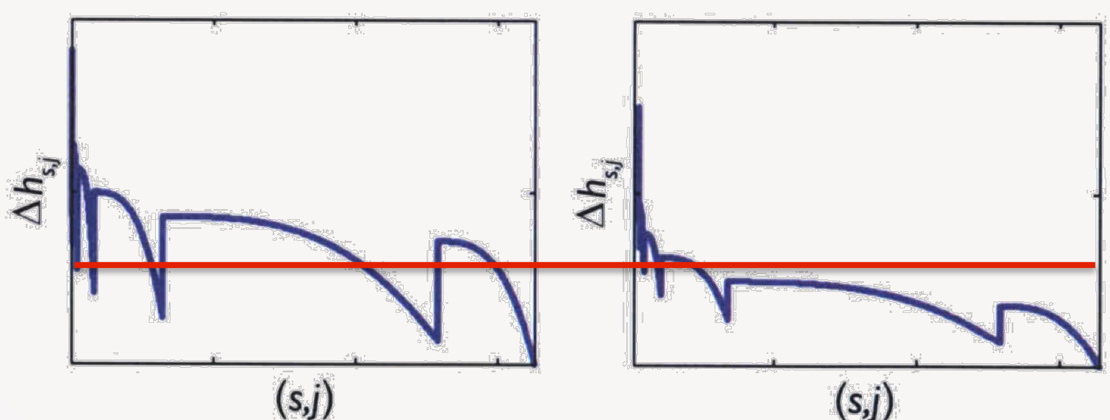

Net Capacity

Experimental Results

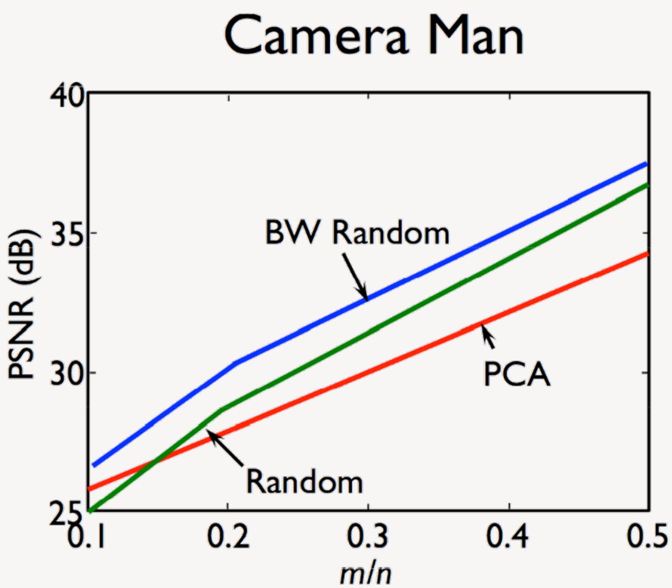

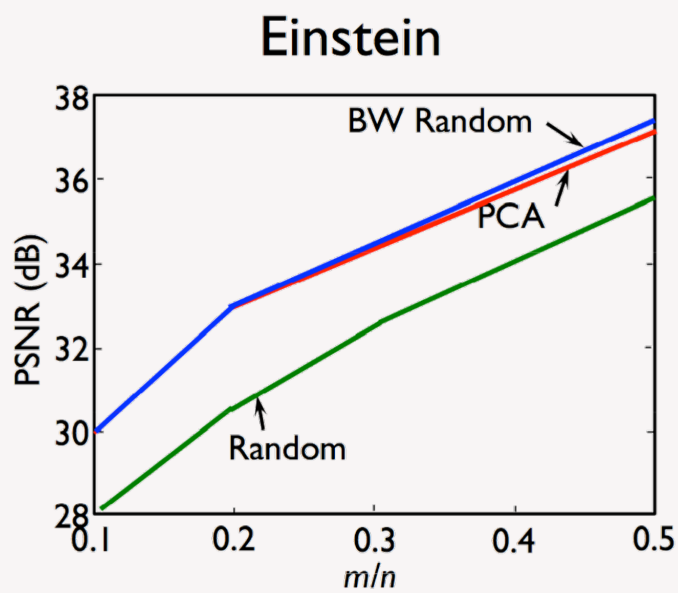

Experimental Results

n=65,536, m=5,000 (7%)

PCA
24.81dB

BW random
25.73dB

Random
23.78dB

Slide modified from Yair Weiss's talk

General Results

- More images from Berkeley dataset

(0.90)

(0.90)

(1.15)

(0.81)

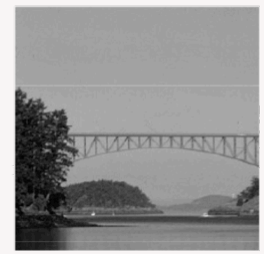

(0.86)

(0.65)

(0.54)

(0.31)

(0.36)

(0.35)

$J_x \approx 0.65$ on average

Slide modified from Yair Weiss's talk

Informative Sensing and Natural Images

25% PCA measurements
(low frequencies, w/ TV min.)
PSNR = 35.96 dB

CS: 25% random measurements

PSNR = 28.02 dB

Slide modified from Yair Weiss's talk

Summary for the CS / IS Theory

- CS theorems hold for ideal sparse signals, but *not* natural images
- Informative sensing aims to design optimal *measurement* matrix given knowledge of signal statistics
- For natural images, optimal measurements depend on *sparsity*
  - For most cases, measuring *low frequencies* is optimal
  - For highly sparse cases, *random projections* are optimal
- Even with optimal projections, the compression gains are modest

Slide Credit: Yair Weiss

Part III:

CS Applications in Vision & Graphics

Beyond Image Reconstruction…

- CS have been adopted to work on many problems in computer vision or computer graphics
- A more useful framework inspired by CS is *sparse representation* and *low-rank approaches*
- Some applications
  - Image Denoising [M. Elad CVPR'06]
  - Face Recognition [J. Wright PAMI'09, A. Wagner PAMI'12]
  - Background subtraction [V. Cevher ECCV'08]
  - Tracking [H. Li CVPR'11, K. Zhang ECCV'12]
  - Media Recovery [J. Gu ECCV'08, PAMI'13]
  - …

CS-inspired Frameworks

- Target for Reconstruction
  - General images or Specific Objects?
- Sparsity Verification
  - Explore a basis to decompose the target into a sparse signal
  - Instead of generic basis like DCT or Wavelet, overcomplete dictionaries are useful in practice
- Requirement for Computational Resources
  - L1-minimization is generally time-consuming!

CS-inspired Face Recognition

- Performance at a glance
  - CS theory implies that precise *choice of feature* space is not so critical as before; *random features* also work!
  - *Robustness to Occlusion:* Source-and-Error Separation

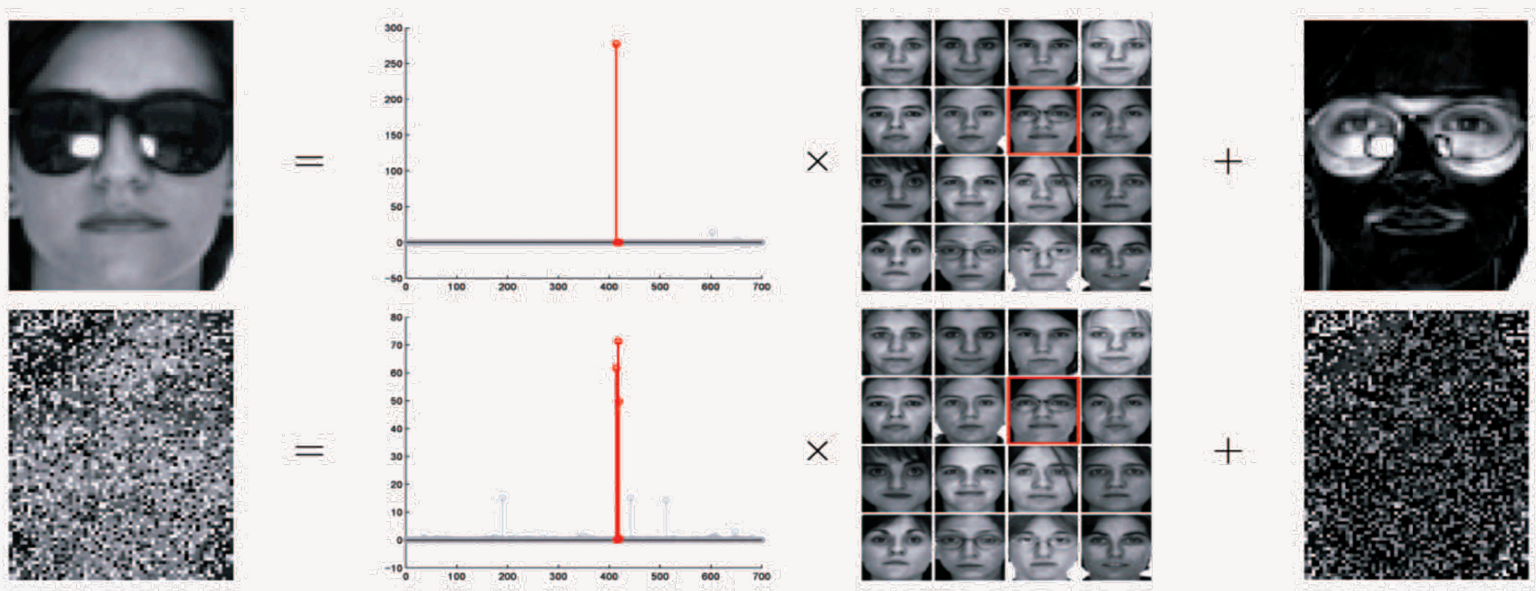

Algorithm

Algorithm 1. Sparse Representation-based Classification (SRC)

1: **Input:** a matrix of training samples $A = [A_1, A_2, \ldots, A_k] \in \mathbb{R}^{m \times n}$ for k classes, a test sample $\boldsymbol{y} \in \mathbb{R}^m$, (and an optional error tolerance $\varepsilon > 0$.)

2: Normalize the columns of A to have unit ℓ^2-norm.

3: Solve the ℓ^1-minimization problem:

$$\hat{\boldsymbol{x}}_1 = \arg\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_1 \quad \text{subject to} \quad A\boldsymbol{x} = \boldsymbol{y}. \tag{13}$$

(Or alternatively, solve

$$\hat{\boldsymbol{x}}_1 = \arg\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_1 \quad \text{subject to} \quad \|A\boldsymbol{x} - \boldsymbol{y}\|_2 \leq \varepsilon.)$$

4: Compute the residuals $r_i(\boldsymbol{y}) = \|\boldsymbol{y} - A\,\delta_i(\hat{\boldsymbol{x}}_1)\|_2$ for $i = 1, \ldots, k$.

5: **Output:** $\text{identity}(\boldsymbol{y}) = \arg\min_i r_i(\boldsymbol{y})$.

Robust Classification – SCI

- Sparsity Concentration Index (SCI)

Projected Sparse Coefficients

$$\mathrm{SCI}(\boldsymbol{x}) \doteq \frac{k \cdot \max_i \|\delta_i(\boldsymbol{x})\|_1 / \|\boldsymbol{x}\|_1 - 1}{k - 1} \quad \in [0, 1]$$

  - SCI = 1, using only images from a single object, valid test image
  - SCI = 0, the sparse coefficients are spread evenly over all classes, *invalid* test image
  - Thresholding on SCI for validation

Sparse Representation Classification (SRC)

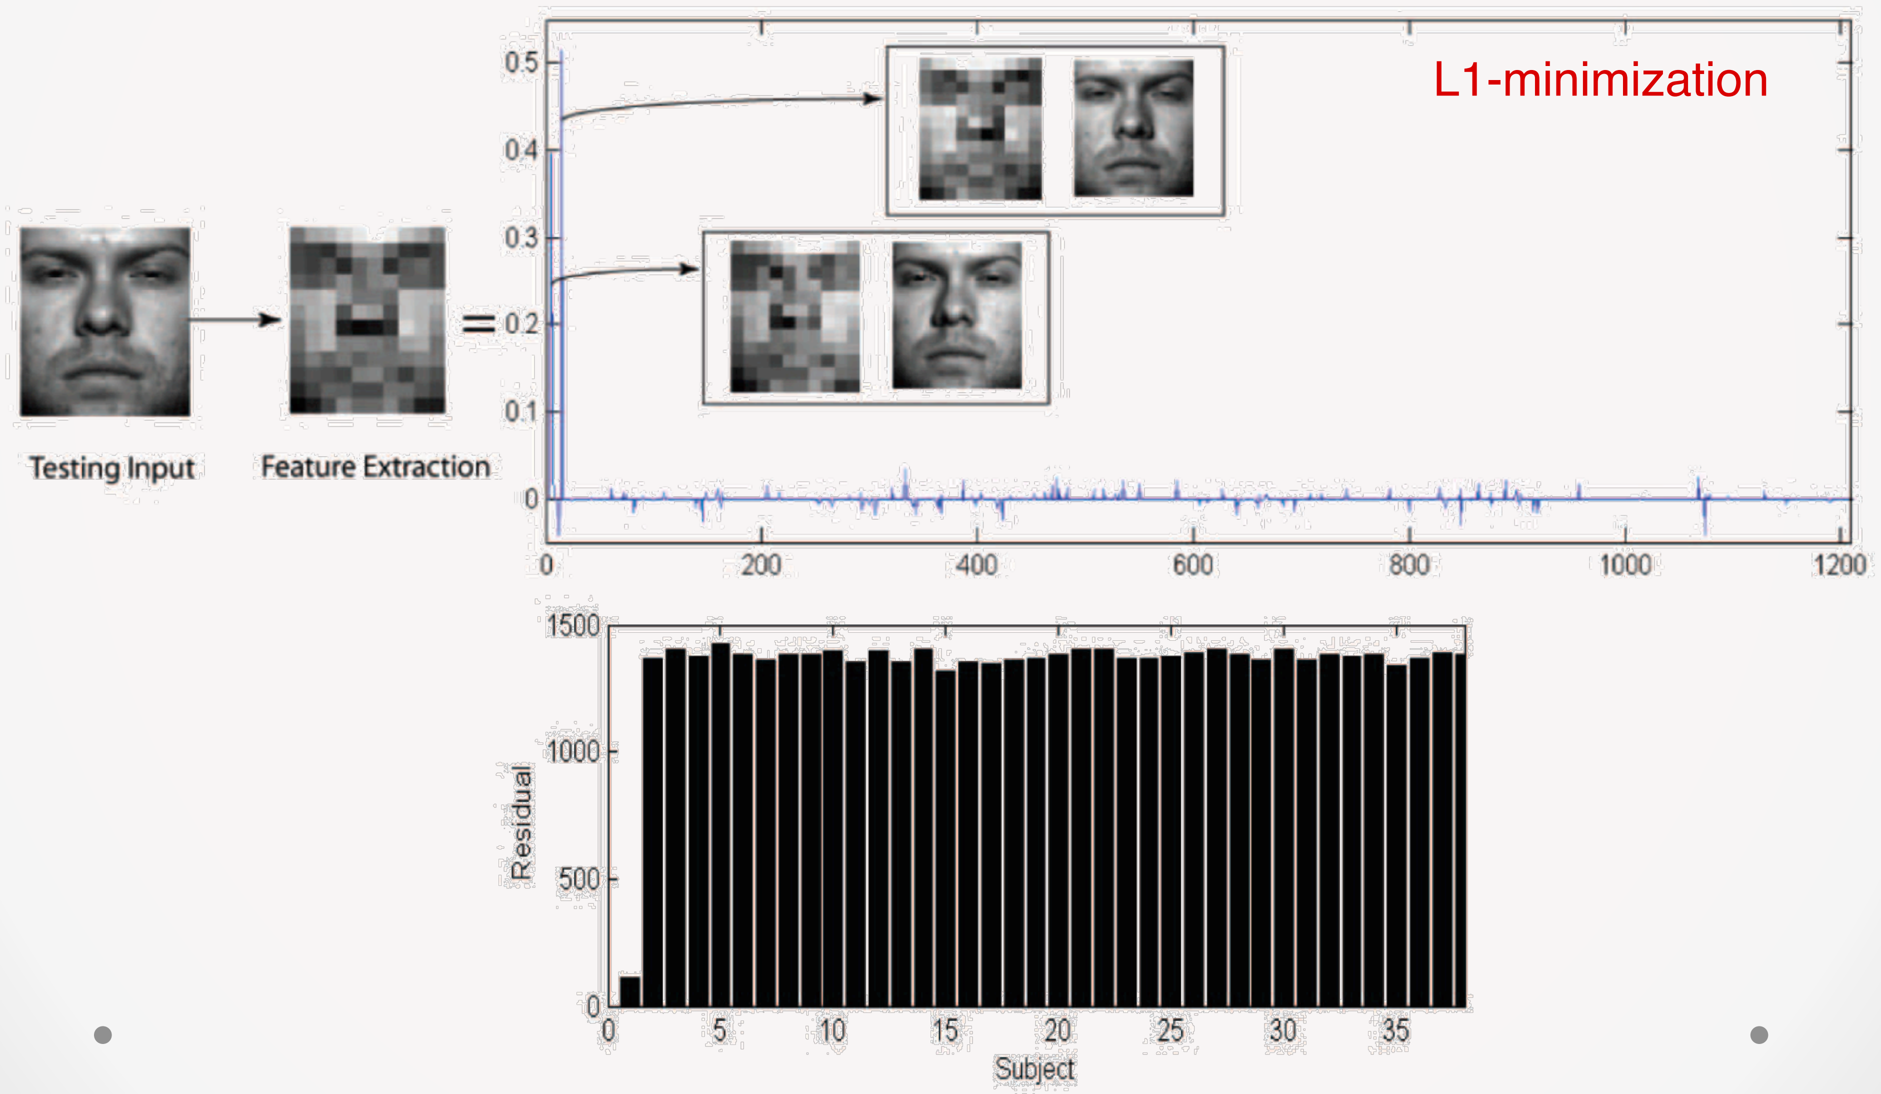

Sparse Representation Classification (SRC)

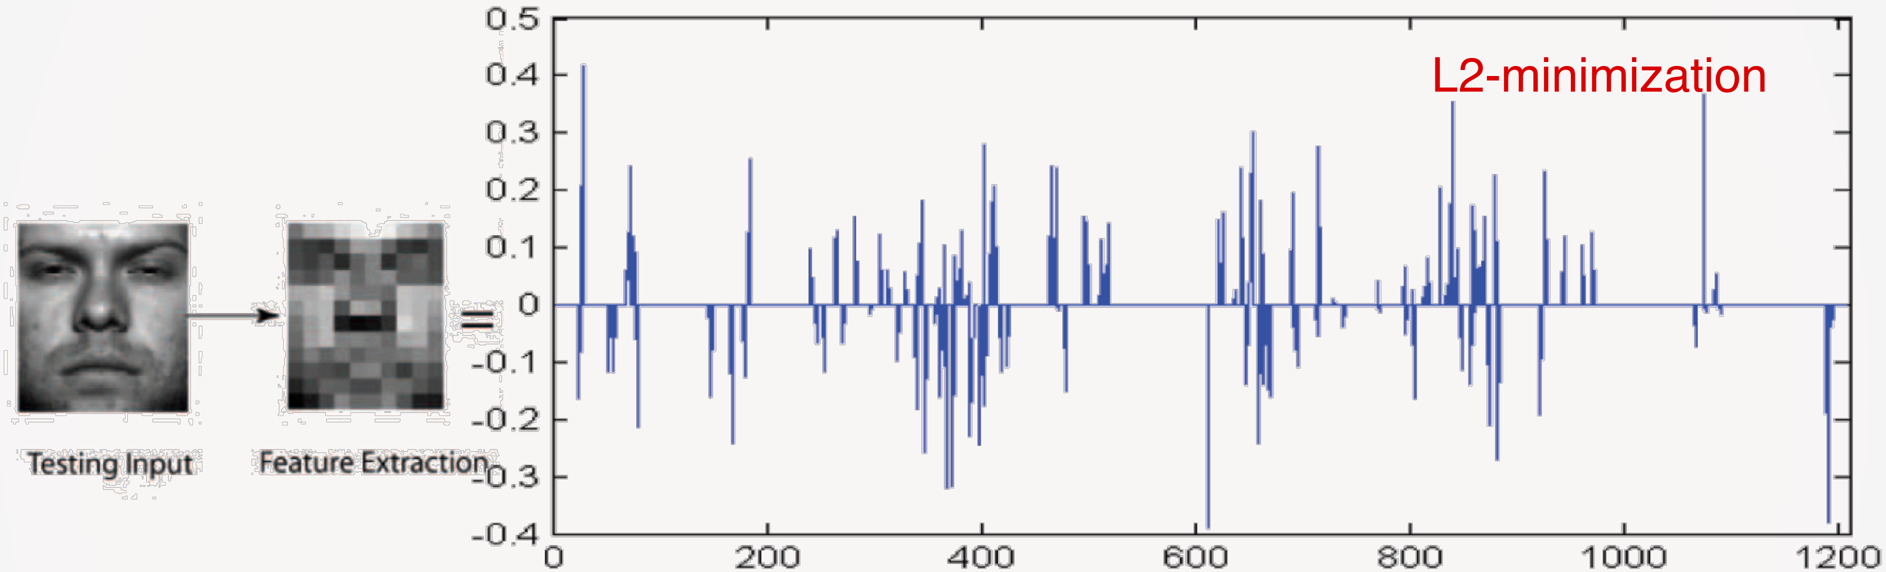

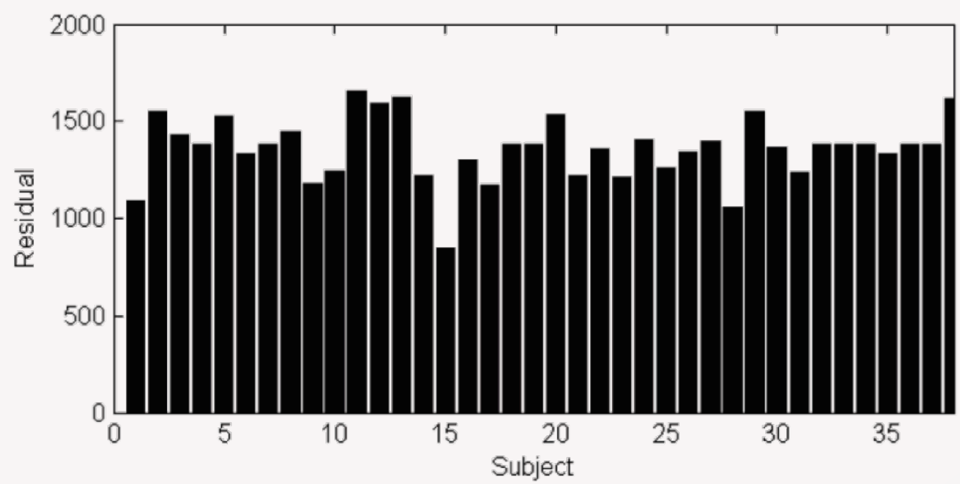

SCI Validation

Invalid Test Image

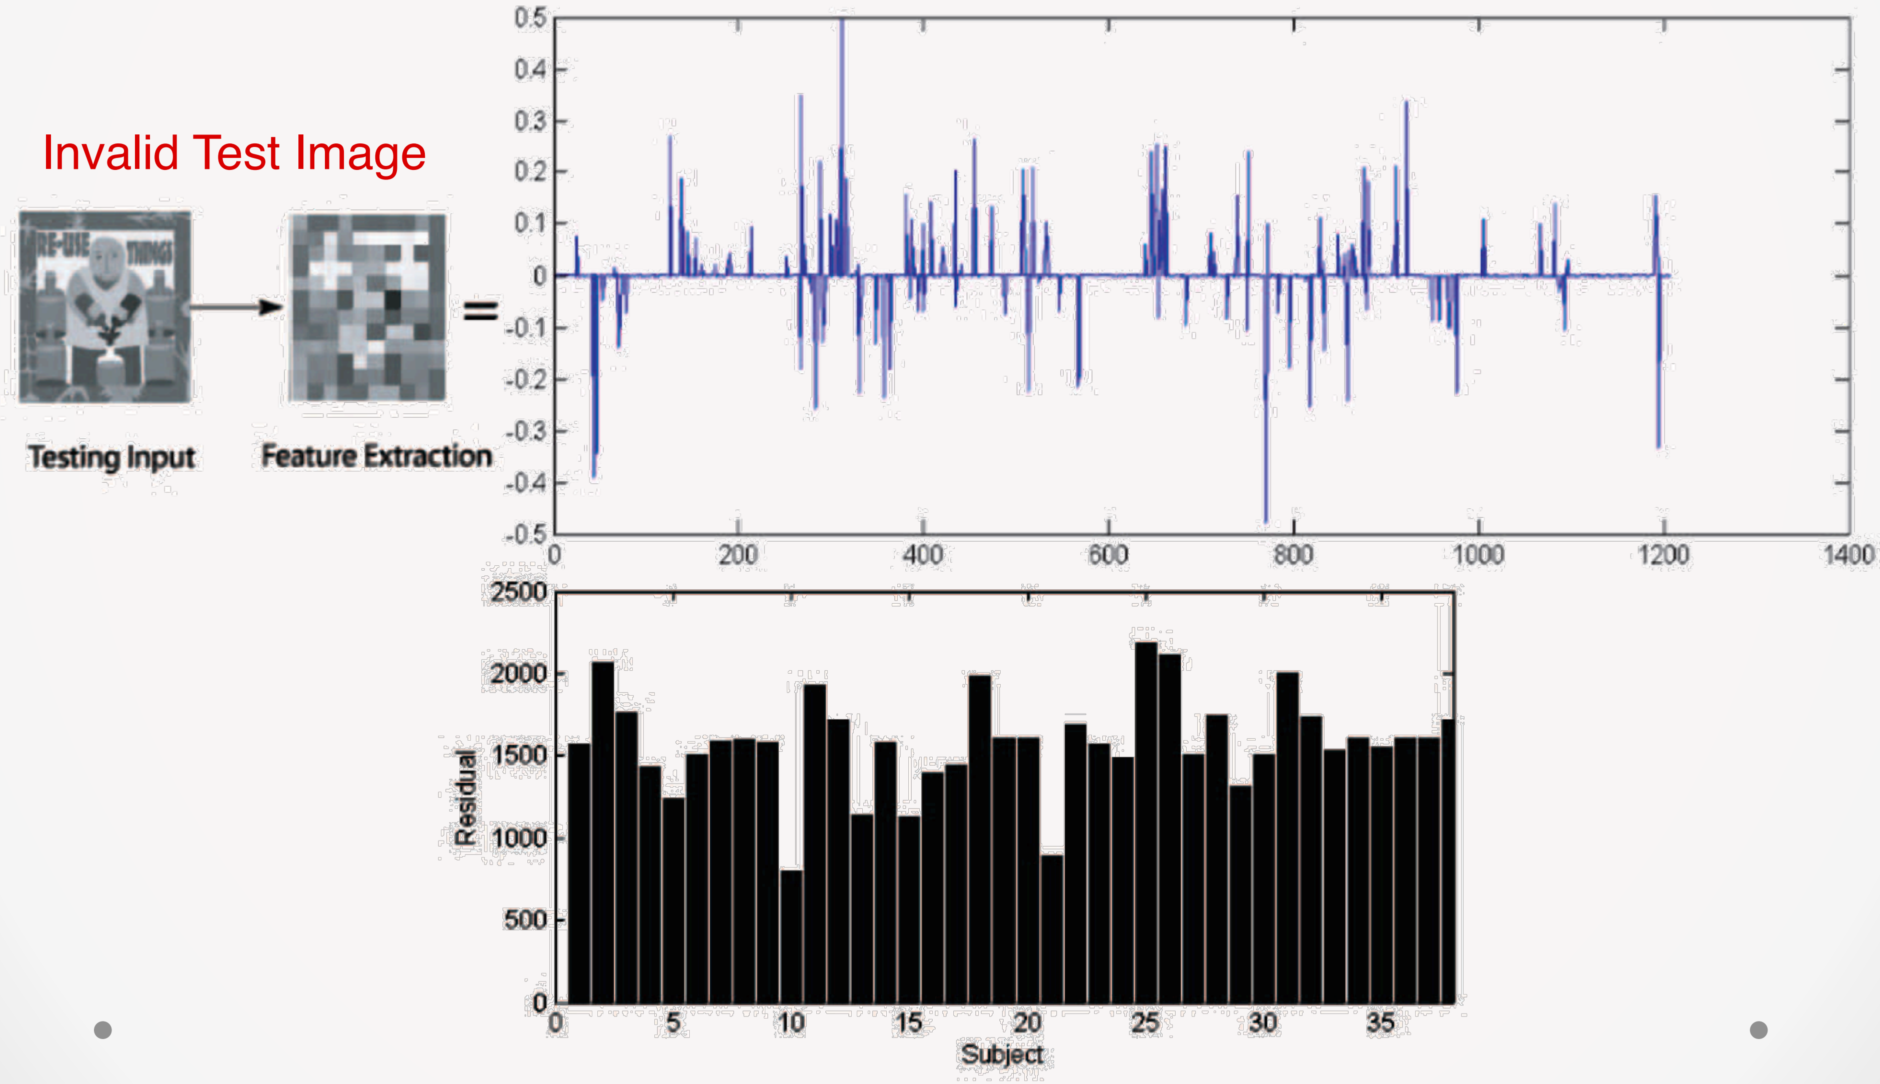

Robust Classification under Occlusion

- Main Concept: occlusions are sparse in some domains

$$y = [A,\ I]\begin{bmatrix} x_0 \\ e_0 \end{bmatrix} \doteq B\,w_0.$$

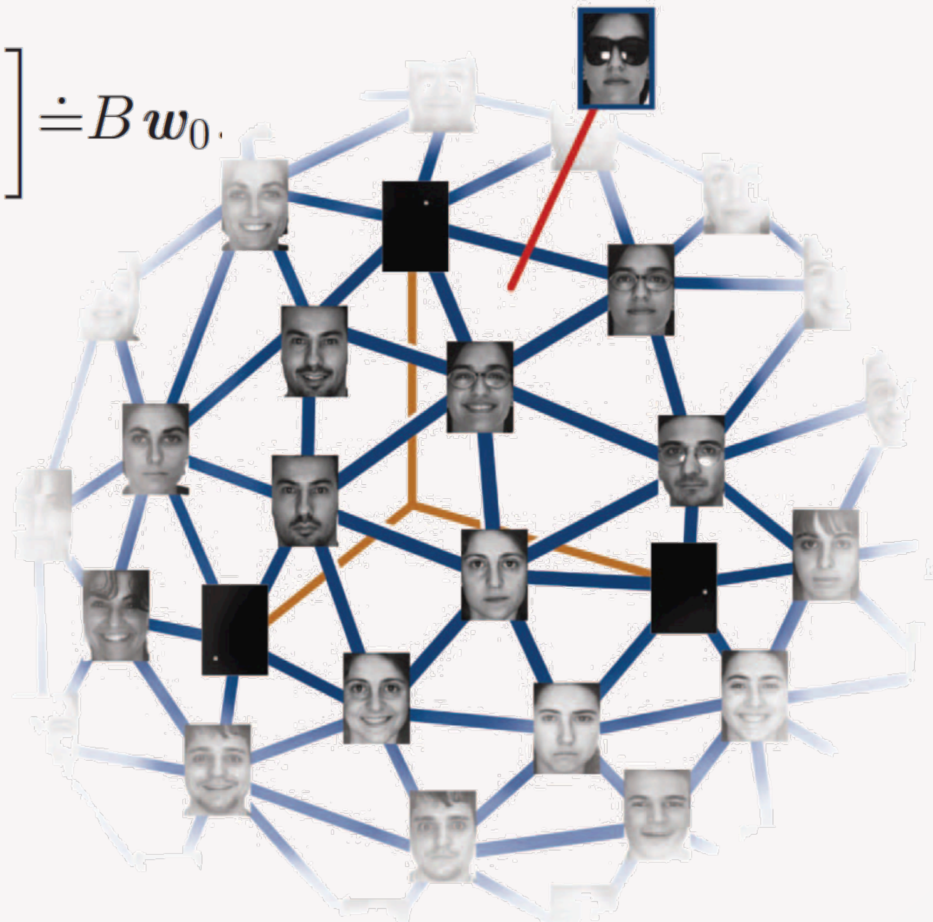

Robust Classification under Occlusion

- Modifications for the Optimization Problem

$$\boldsymbol{y} = B\boldsymbol{w} \quad \text{with} \quad B = [A, \; A_e] \quad \in \mathbb{R}^{m \times (n+n_e)}$$

Sparse Basis for Noise

$$\hat{\boldsymbol{w}}_1 = \arg\min \|\boldsymbol{w}\|_1 \quad \text{subject to} \quad B\boldsymbol{w} = \boldsymbol{y}$$

  - If the proportion of occlusion is not too large, it is still possible to reconstruct the original sparse representation
  - Ideally, we want (verified by experiments)

$$n_i + |\text{support}(\boldsymbol{e}_0)| > d/3$$

Sparsity of Original Face

Sparsity of Noise or Occlusion

Original Dimension

Datasets

- Extended Yale B Database
  - 2,414 frontal-face images of 38 individuals
  - 192x168 normalized face images
  - Captured under various laboratory-controlled lighting conditions
- AR Database
  - Over 4,000 frontal images for 126 individuals
  - Include more facial variations, including illumination change, expressions, and facial disguises
  - Select 50 male and female subjects with 14 variations
  - Normalize image size to 165x120

Feature Extraction

$\overset{R}{\Longrightarrow}$

Eigenface

Laplacianface

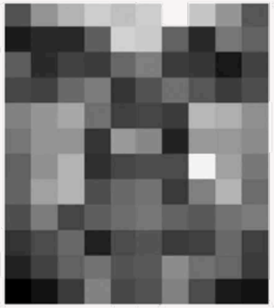

Downsample

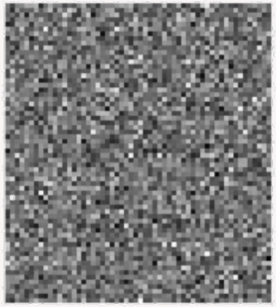

Random Projection

Extended Yale B Database

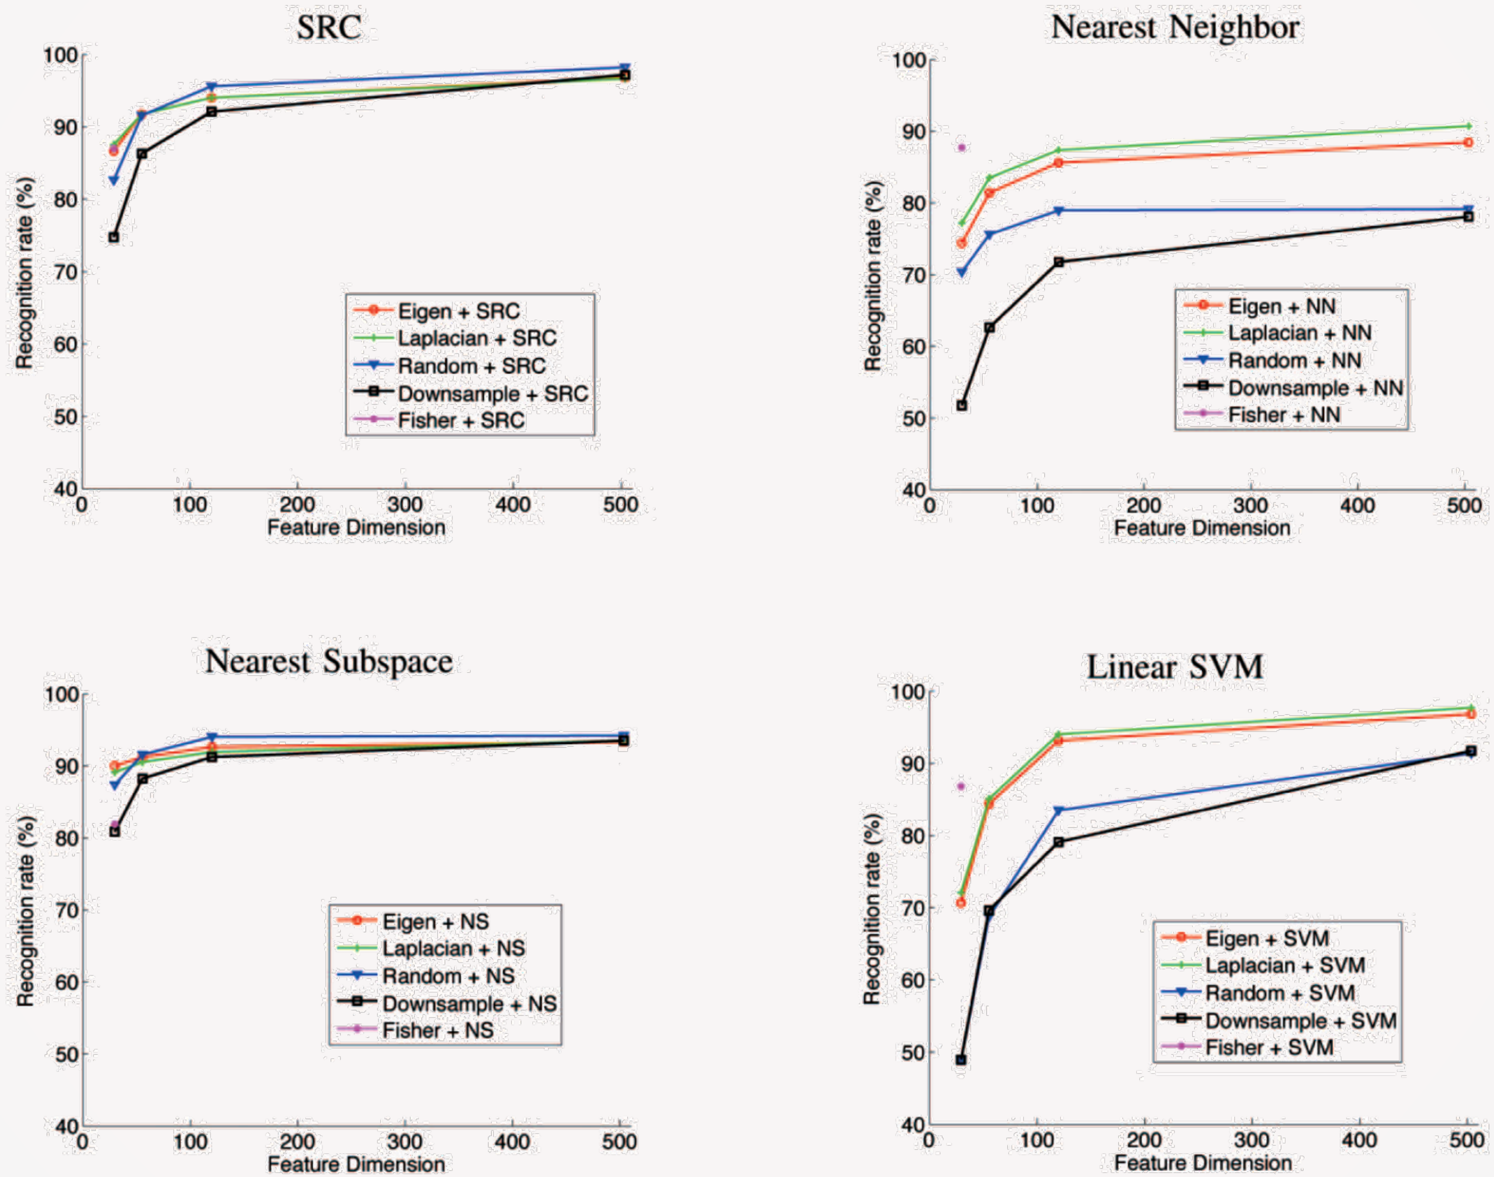

AR Database

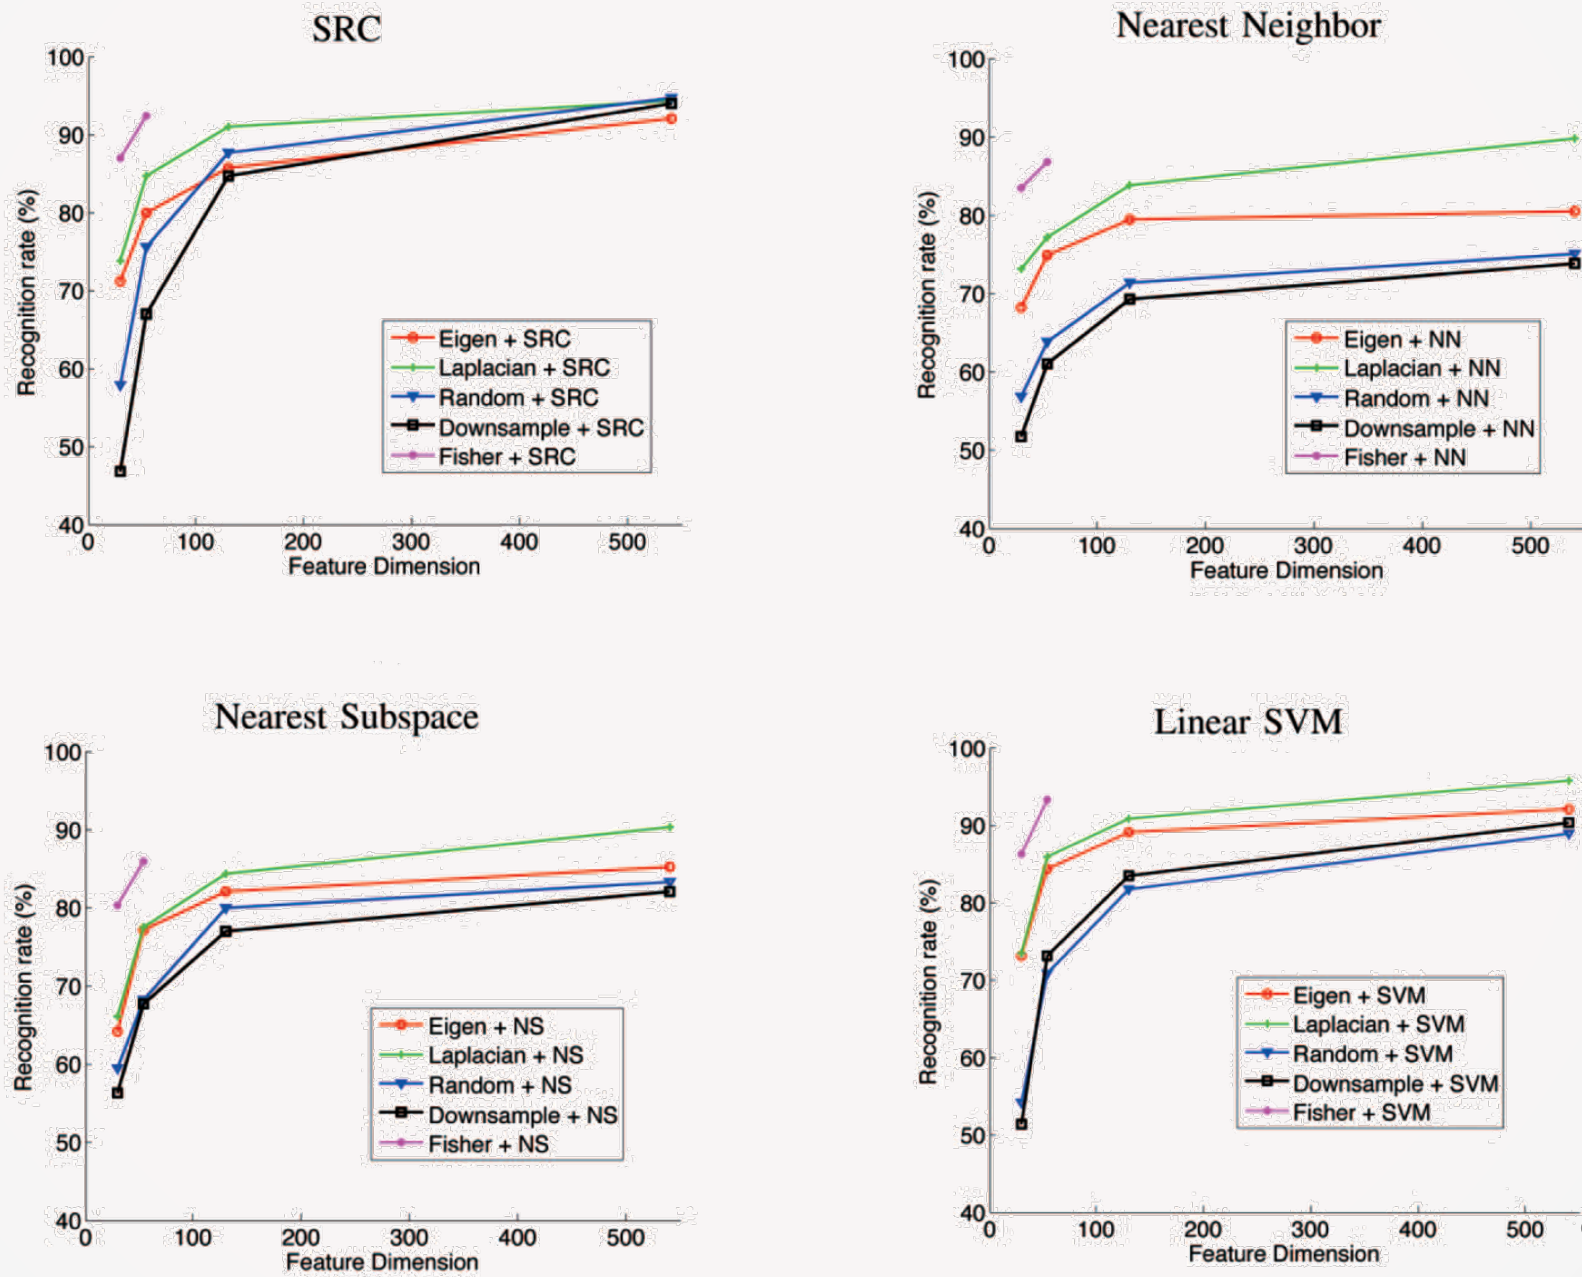
SRC
Recognition rate (%)
Feature Dimension
Eigen + SRC
Laplacian + SRC
Random + SRC
Downsample + SRC
Fisher + SRC
Nearest Neighbor
Recognition rate (%)
Feature Dimension
Eigen + NN
Laplacian + NN
Random + NN
Downsample + NN
Fisher + NN
Nearest Subspace
Recognition rate (%)
Feature Dimension
Eigen + NS
Laplacian + NS
Random + NS
Downsample + NS
Fisher + NS
Linear SVM
Recognition rate (%)
Feature Dimension
Eigen + SVM
Laplacian + SVM
Random + SVM
Downsample + SVM
Fisher + SVM

Recognition with Partial Face Features

Features	Nose	Right Eye	Mouth & Chin
Dimension (d)	4,270	5,040	12,936
SRC	**87.3%**	**93.7%**	**98.3%**
NN	49.2%	68.8%	72.7%
NS	83.7%	78.6%	94.4%
SVM	70.8%	85.8%	95.3%

Results under Random Corruption

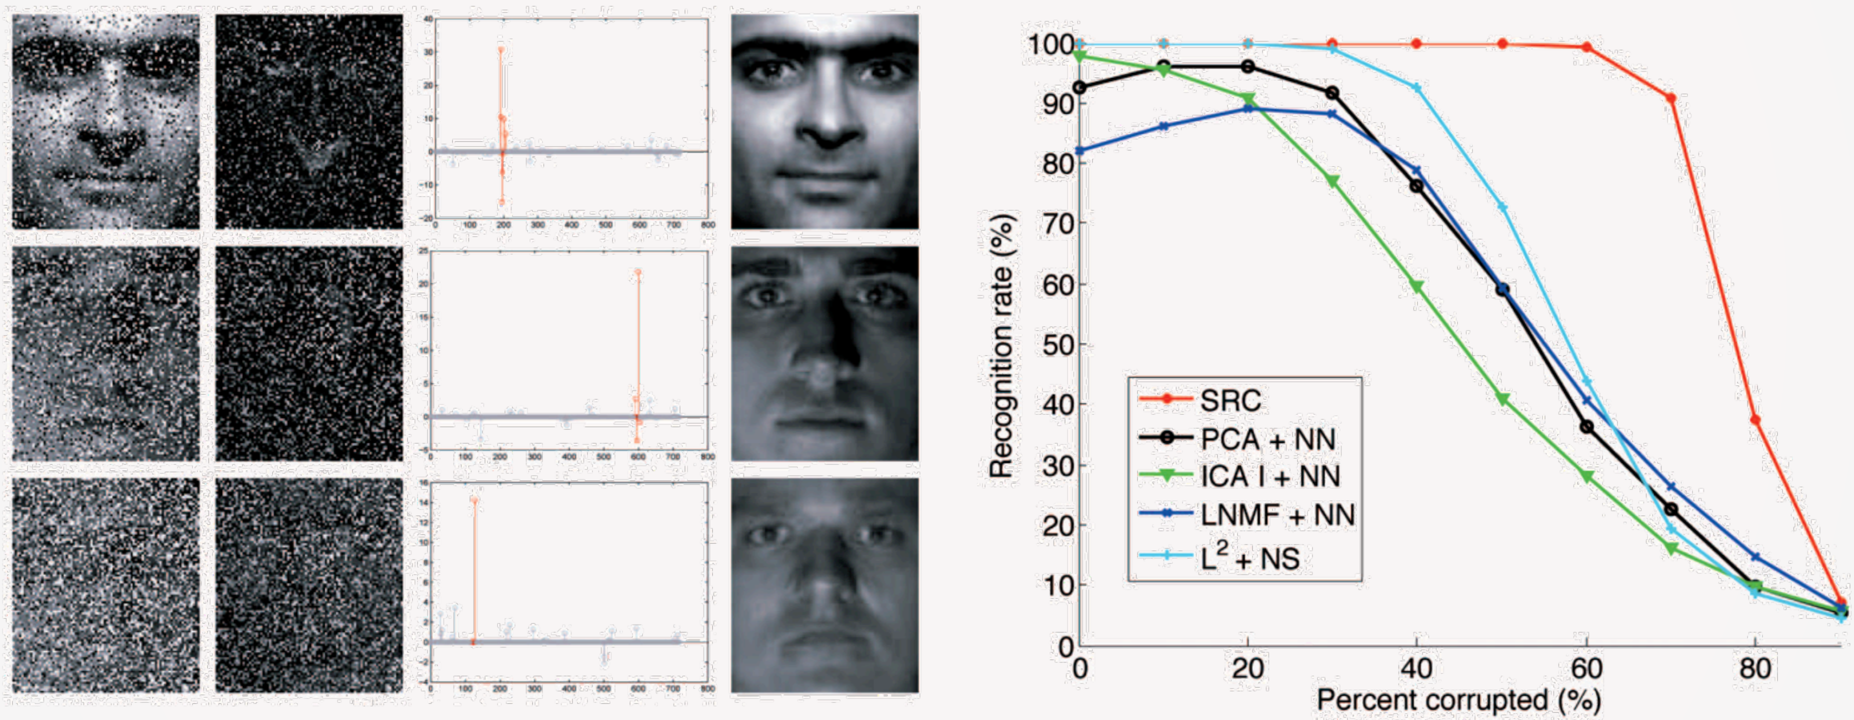

(a) (b) (c) (d) (e)

Percent corrupted	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%
Recognition rate	100%	100%	100%	100%	100%	100%	99.3%	90.7%	37.5%	7.1%

Results under Contiguous Occlusion

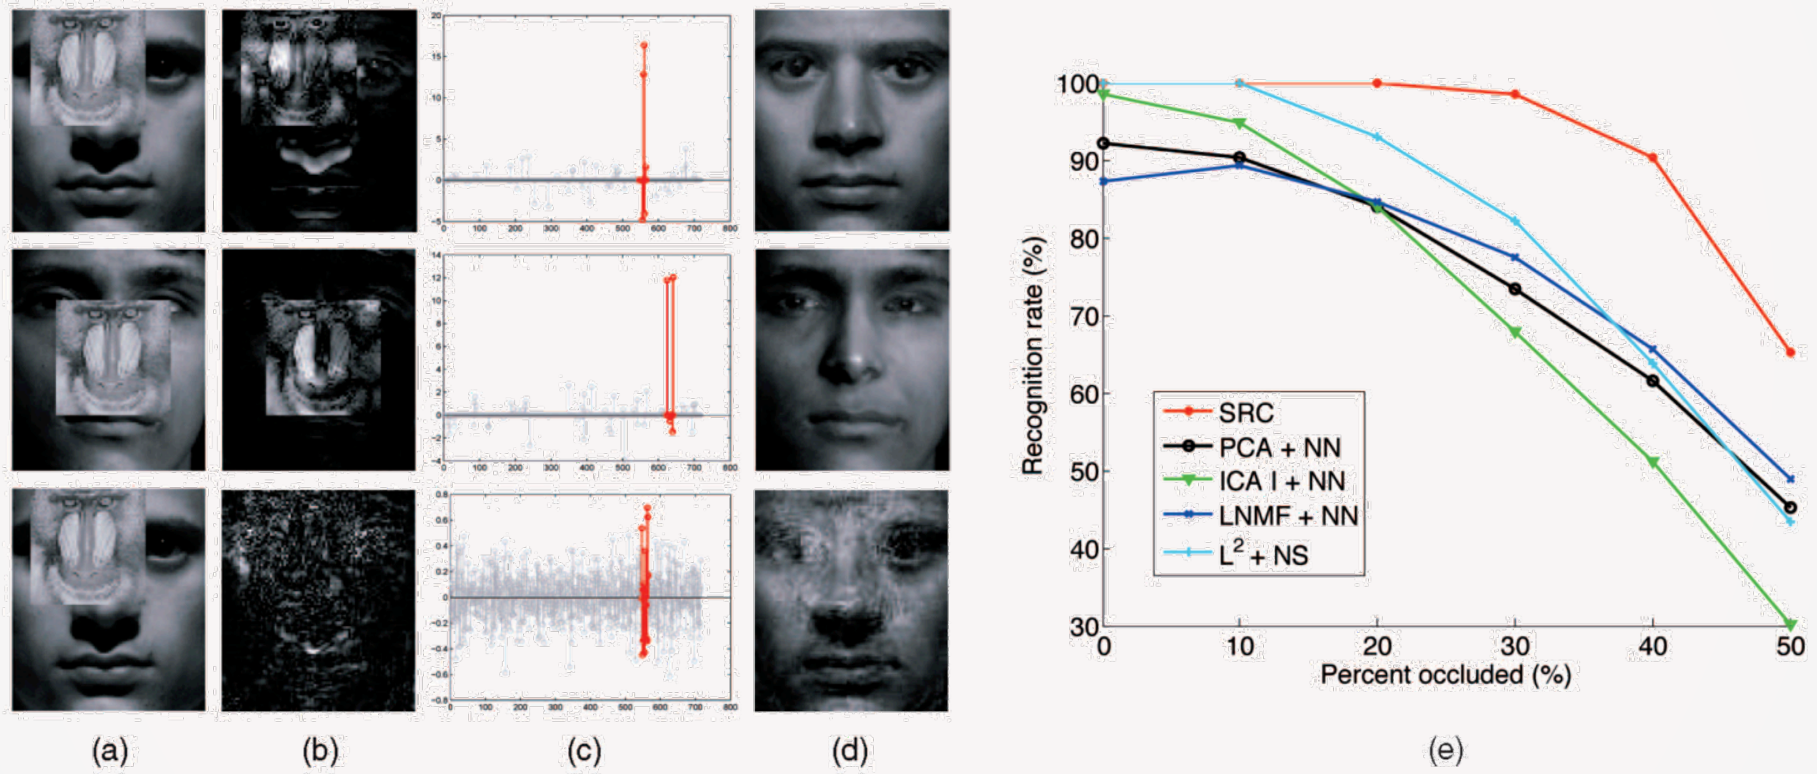

Percent occluded	0%	10%	20%	30%	40%	50%
Recognition rate	100%	100%	99.8%	98.5%	90.3%	65.3%

Partition Scheme

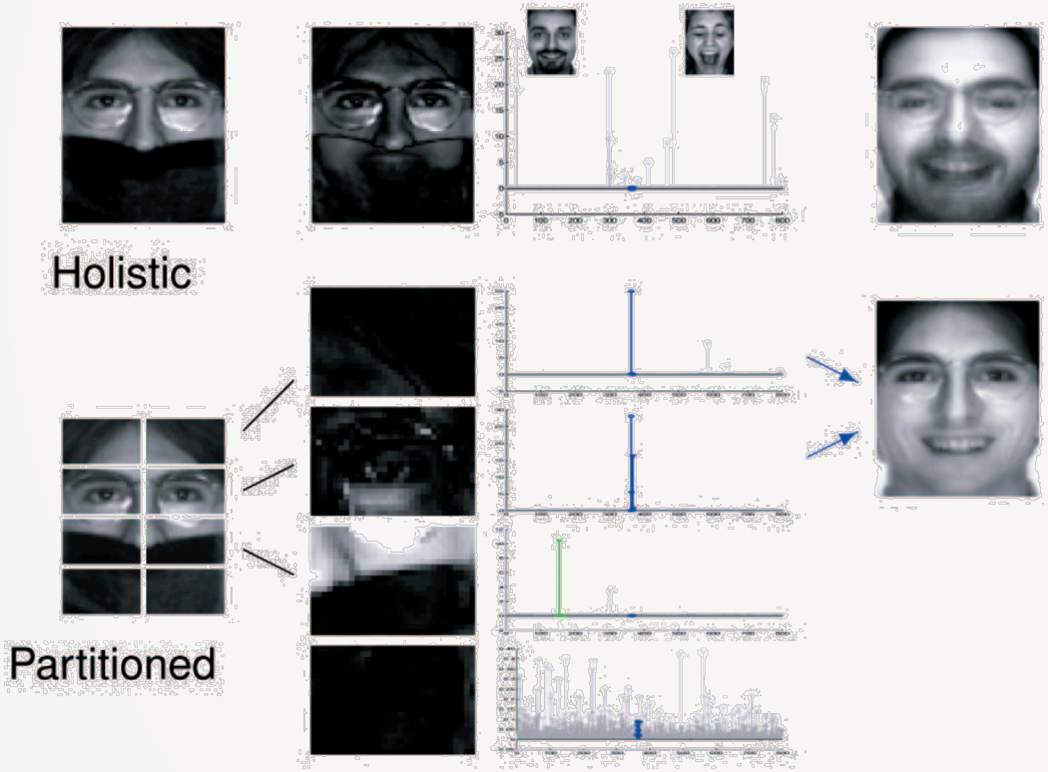

Algorithms	Rec. rate sunglasses	Rec. rate scarves
SRC (partitioned)	**87.0%** **(97.5%)**	**59.5%** **(93.5%)**
PCA + NN	70.0%	12.0%
ICA I + NN	53.5%	15.0%
LNMF + NN	33.5%	24.0%
ℓ^2 + NS	64.5%	12.5%

Summary & Comment

- The CS-inspired SRC framework provides a relatively robust recognition framework
  - Feature Extraction: choices of feature still matter
  - Robustness to Occlusion: use *local* (partitioned scheme) to improve global SRC

- Limitations
  - The algorithm is robust to *small* variations (in pose or displacement), but high-quality detection, cropping and normalization are still desired
  - Illumination and deformable problems have to be solved by additional modeling [PAMI'13]

Extension (ECCV'12 short course):
http://www.eecs.berkeley.edu/~yang/courses/ECCV2012/

What's Participating Media

- *Participating Media*: Real-world phenomena that can only be described by volume densities rather than boundary surface
  - Examples are like smoke, clouds, or dynamic liquid

- Conventional ways for recovering the volume densities of participating media include
  - Laser sheet + high-speed camera + scanning
  - Static laser rays + interpolation

- How to improve *efficiency in acquisition* is critical!

Property of Participating Media

- The volume densities of participating media usually have high sparsity (> 0.95) over time

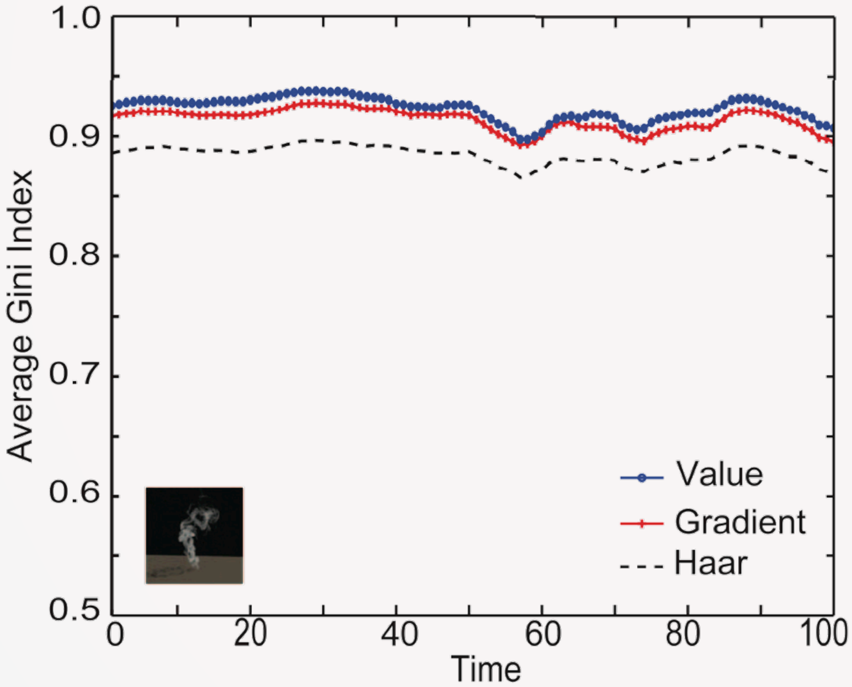

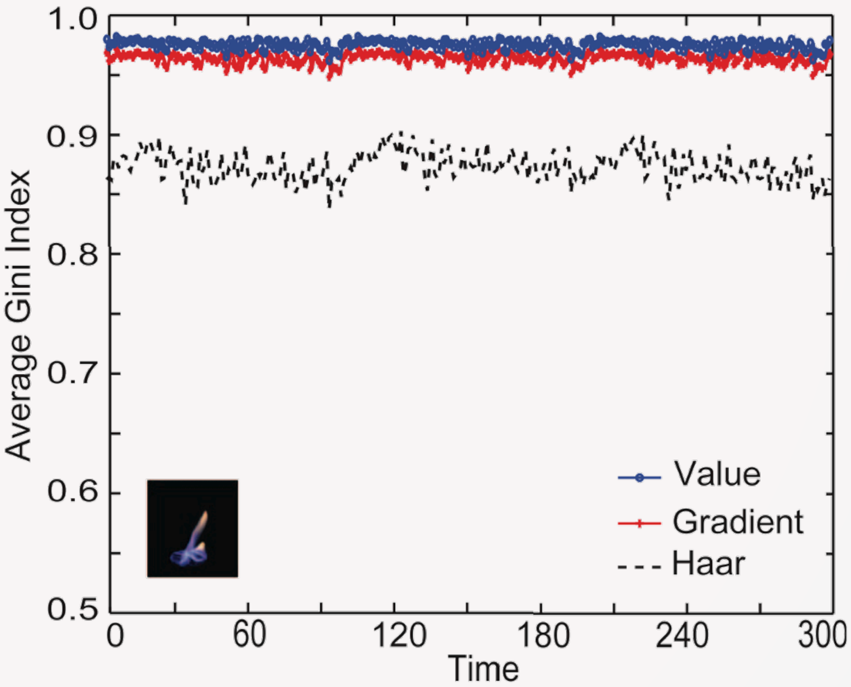

- J. Gu et al, ECCV'08, PAMI'13

Compressive Structured Light

- Settings and Image Formation Model

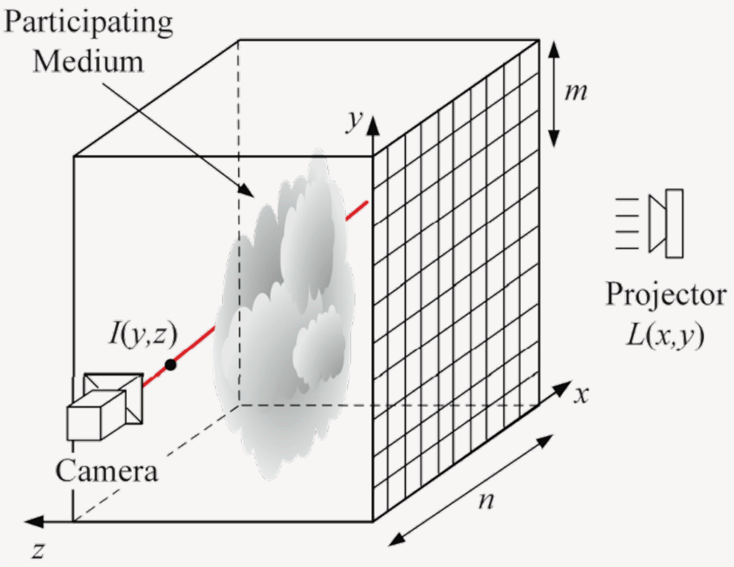

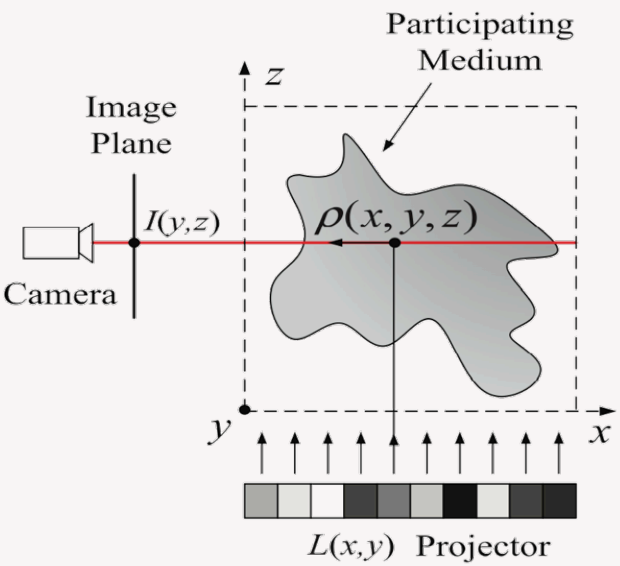

$$I(y, z) = \int_x \rho(x, y, z) \cdot L(x, y)\, dx$$

- J. Gu et al, ECCV'08, PAMI'13

Compressive Structured Light

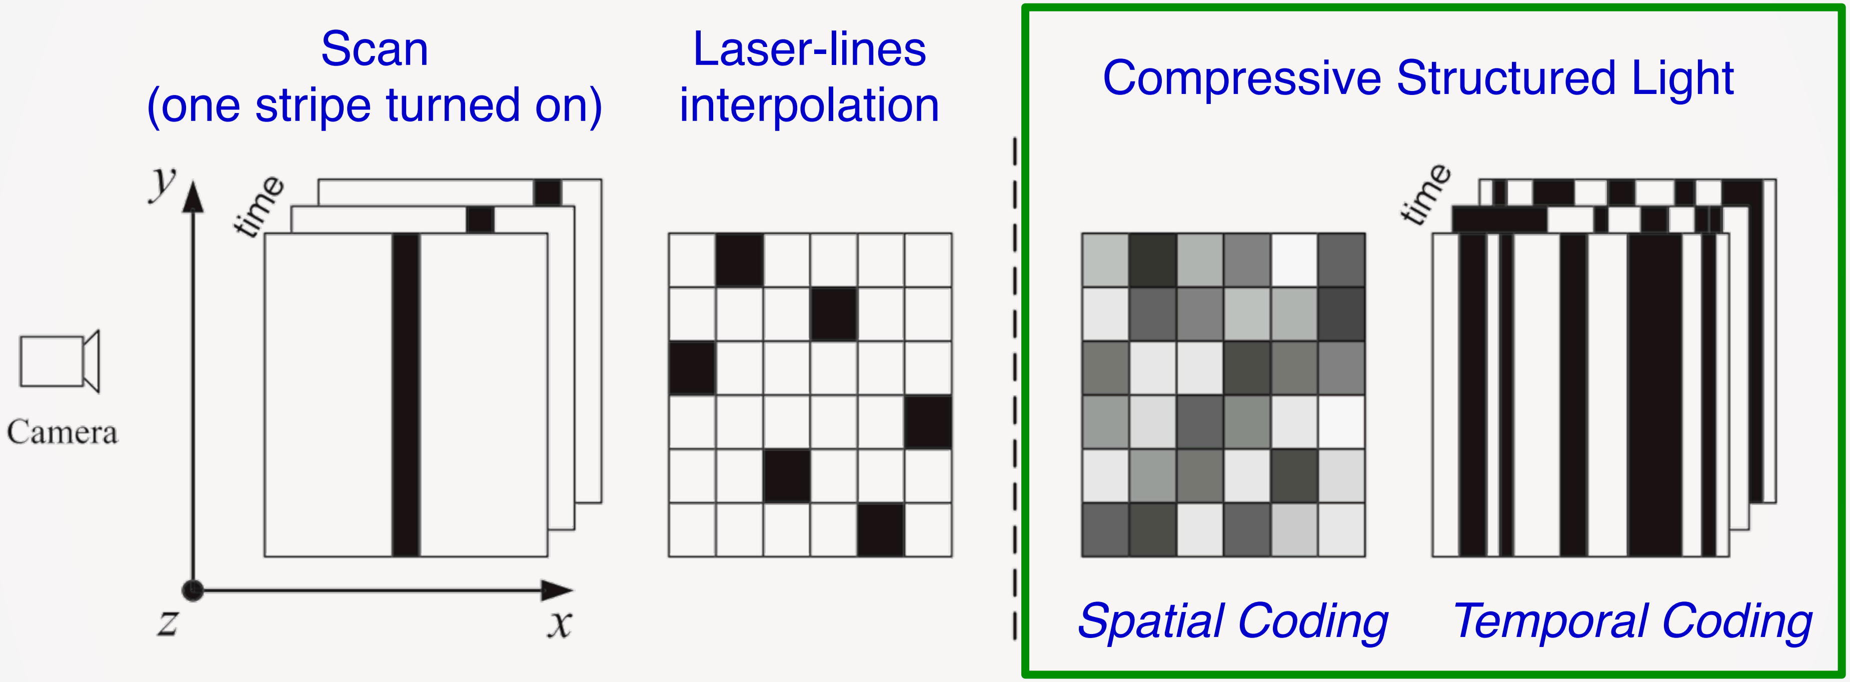

- The utilization of light is more efficient
- The measurement process is highly efficient in both *acquisition time* and *illumination power*

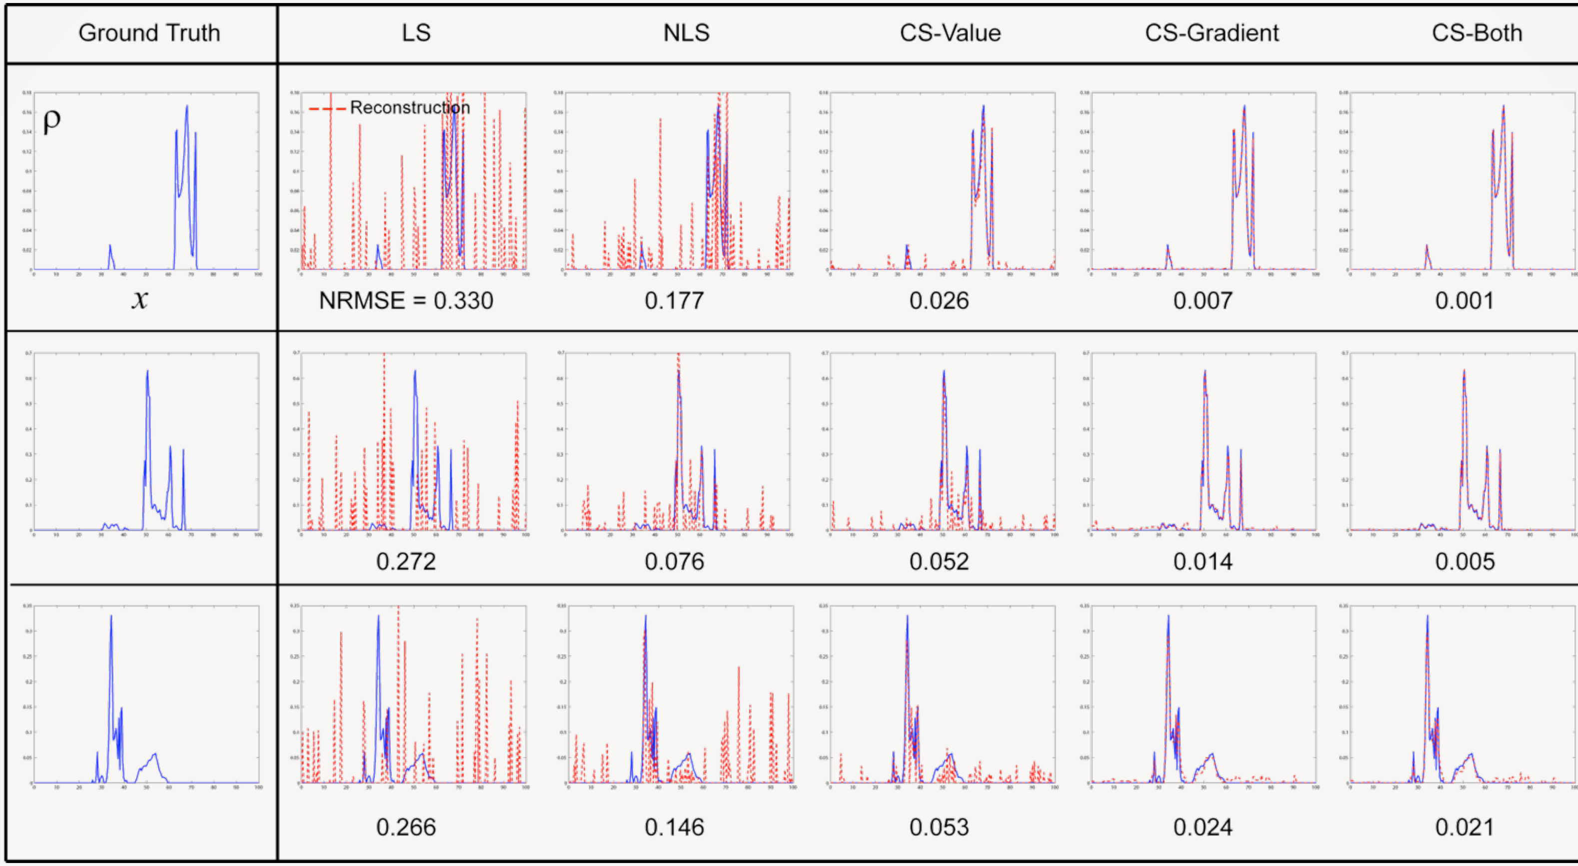

Objective Functionals Used for Volume Reconstruction

Method	Optimization Functional	Constraints
Least Square (LS)	$\|\|\mathbf{Ax}-\mathbf{b}\|\|_2$	
Nonnegative LS	$\|\|\mathbf{Ax}-\mathbf{b}\|\|_2$	$\mathbf{x} \geq 0$
CS-Value	$\|\|\mathbf{x}\|\|_1$	$\mathbf{Ax}=\mathbf{b},\ \mathbf{x} \geq 0$
CS-Gradient	$\|\|\mathbf{x}'\|\|_1$	$\mathbf{Ax}=\mathbf{b},\ \mathbf{x} \geq 0$
CS-Both	$\|\|\mathbf{x}\|\|_1 + \lambda\|\|\mathbf{x}'\|\|_1$	$\mathbf{Ax}=\mathbf{b},\ \mathbf{x} \geq 0$

Simulation Results

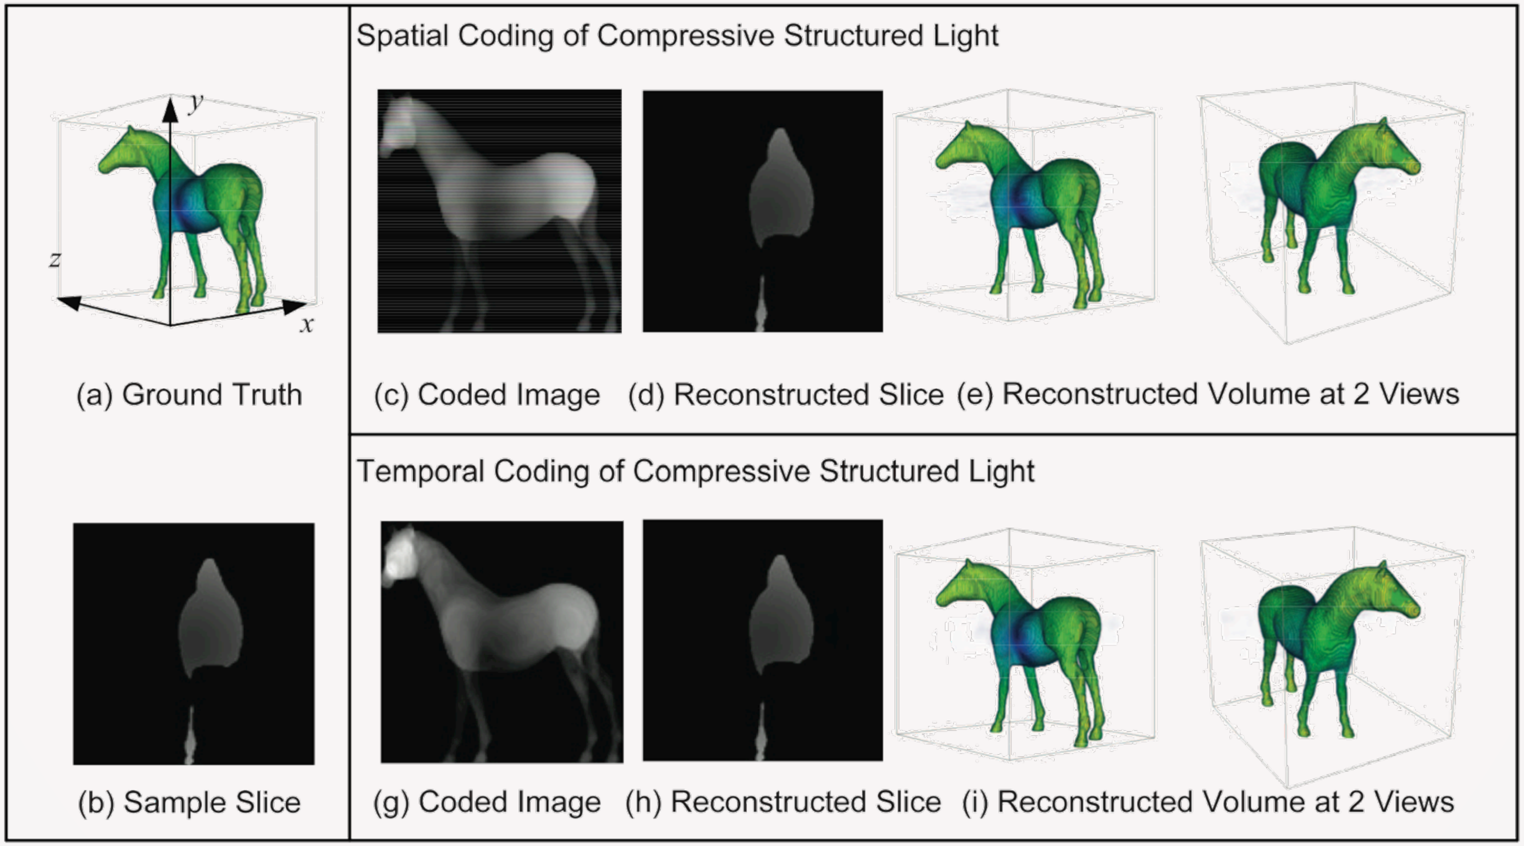

- J. Gu et al, ECCV'08, PAMI'13

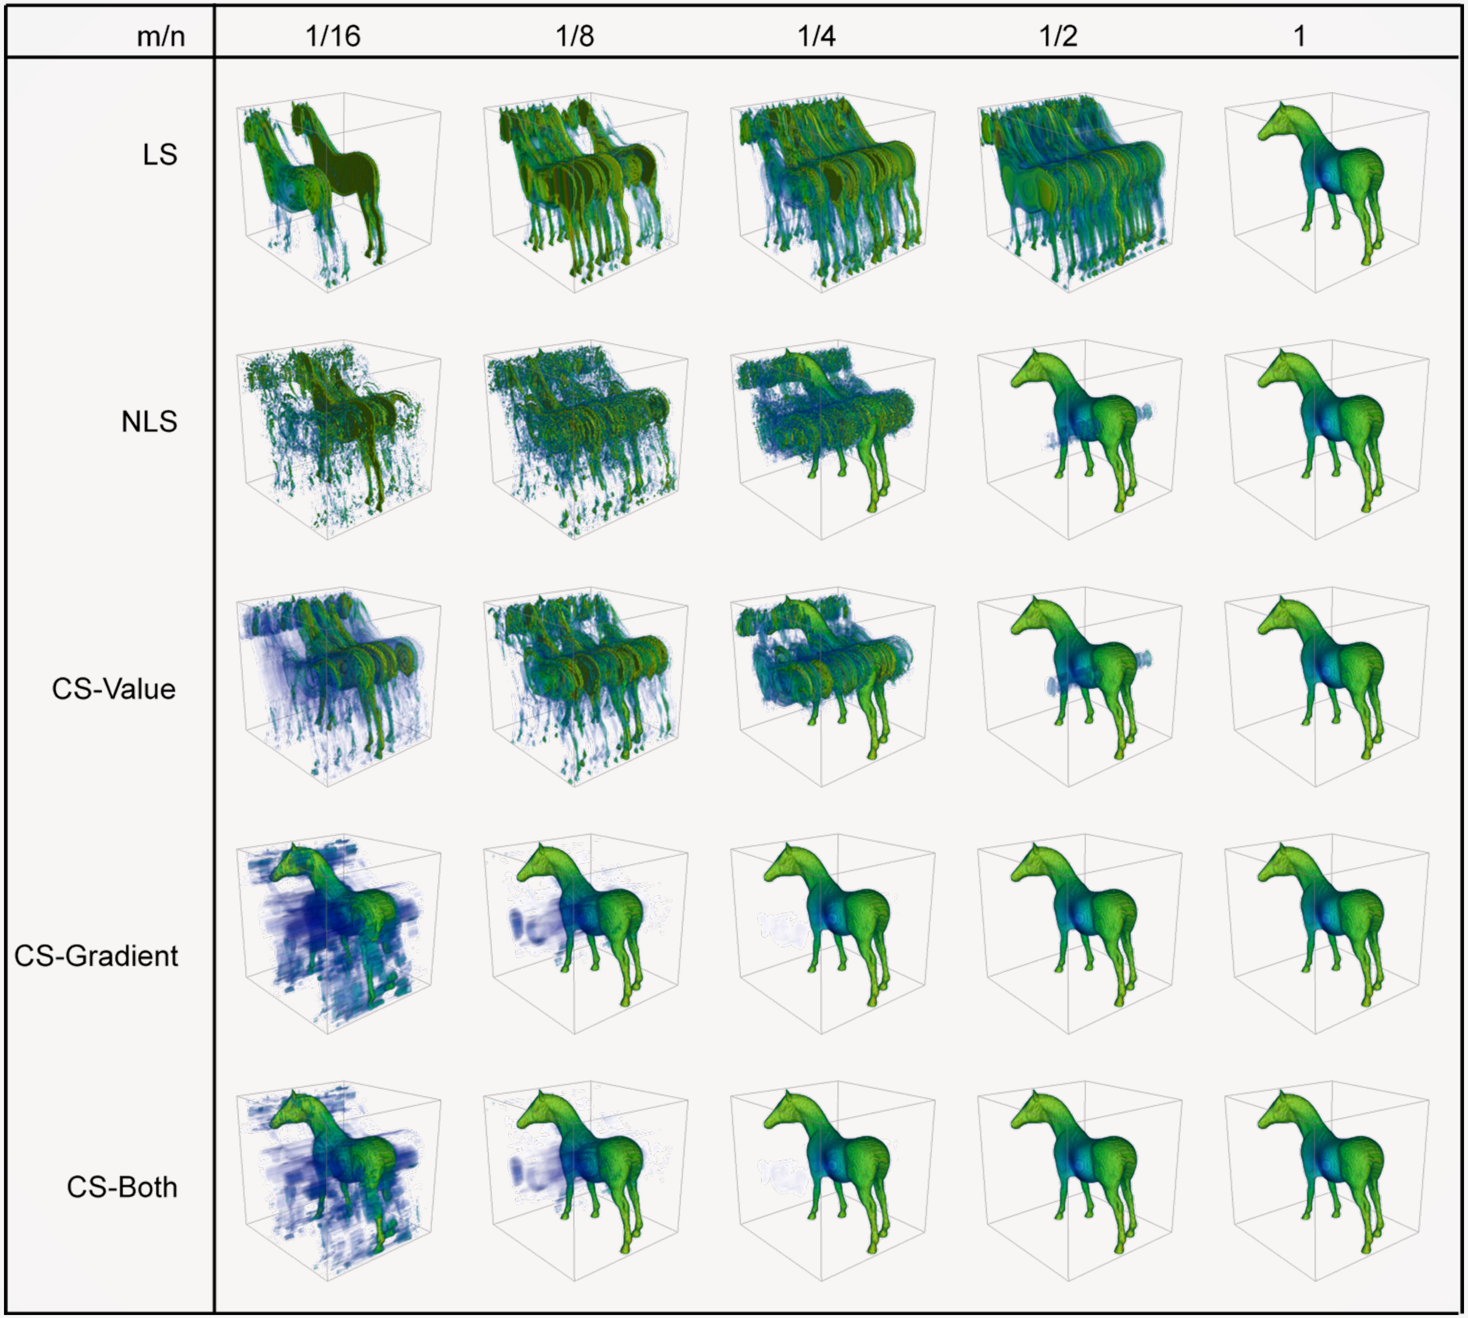

- J. Gu et al, ECCV'08, PAMI'13

Experimental Results – Static Volumes

- Artifacts
  - Multiple Scattering
  - Attenuation

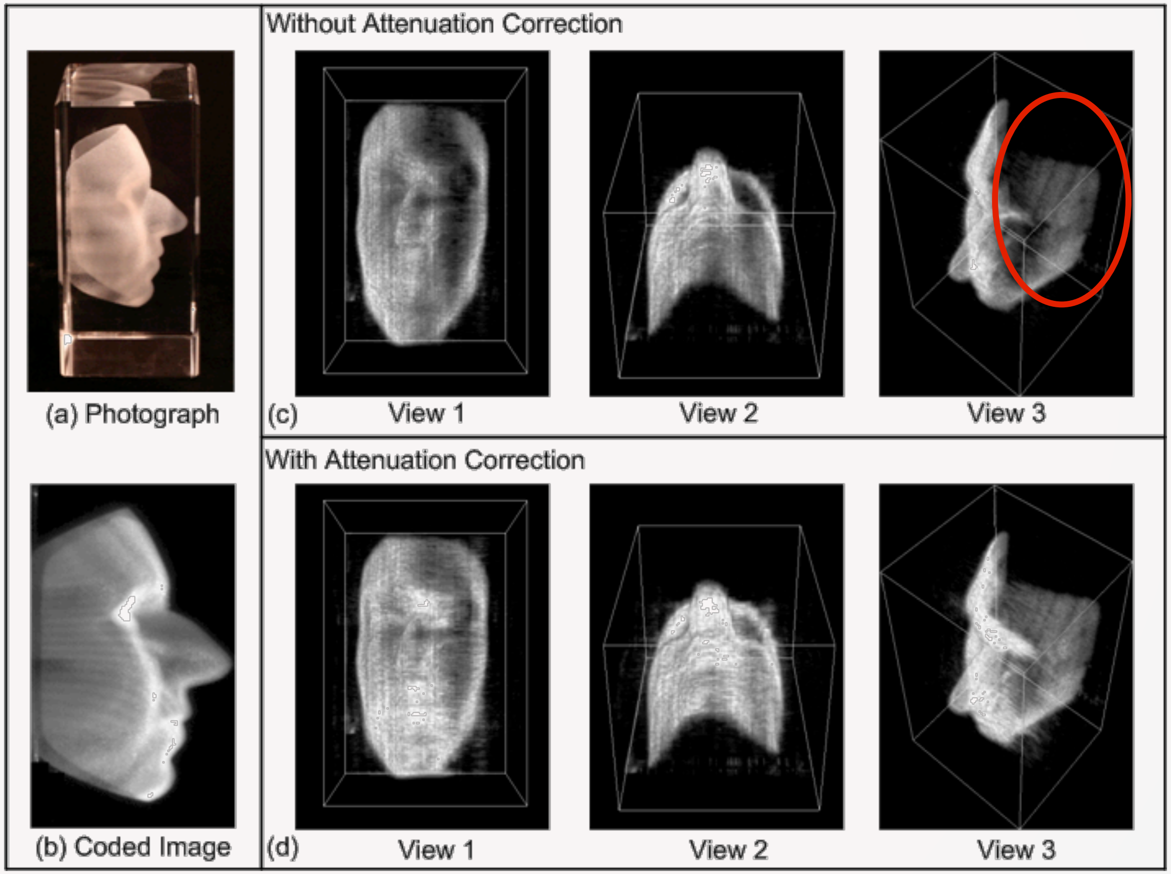

- J. Gu et al, ECCV'08, PAMI'13

Dynamic Volumes – Milk Drop

- Specification: 128 x128 x 128 volume at 15 fps

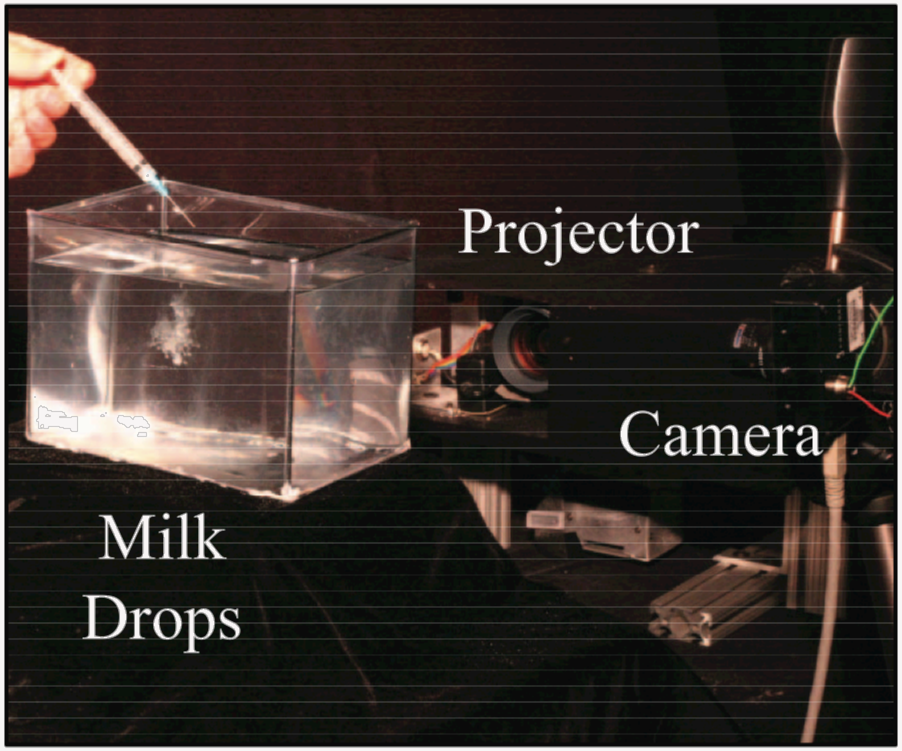

- J. Gu et al, ECCV'08, PAMI'13

Dynamic Volumes – Milk Drop

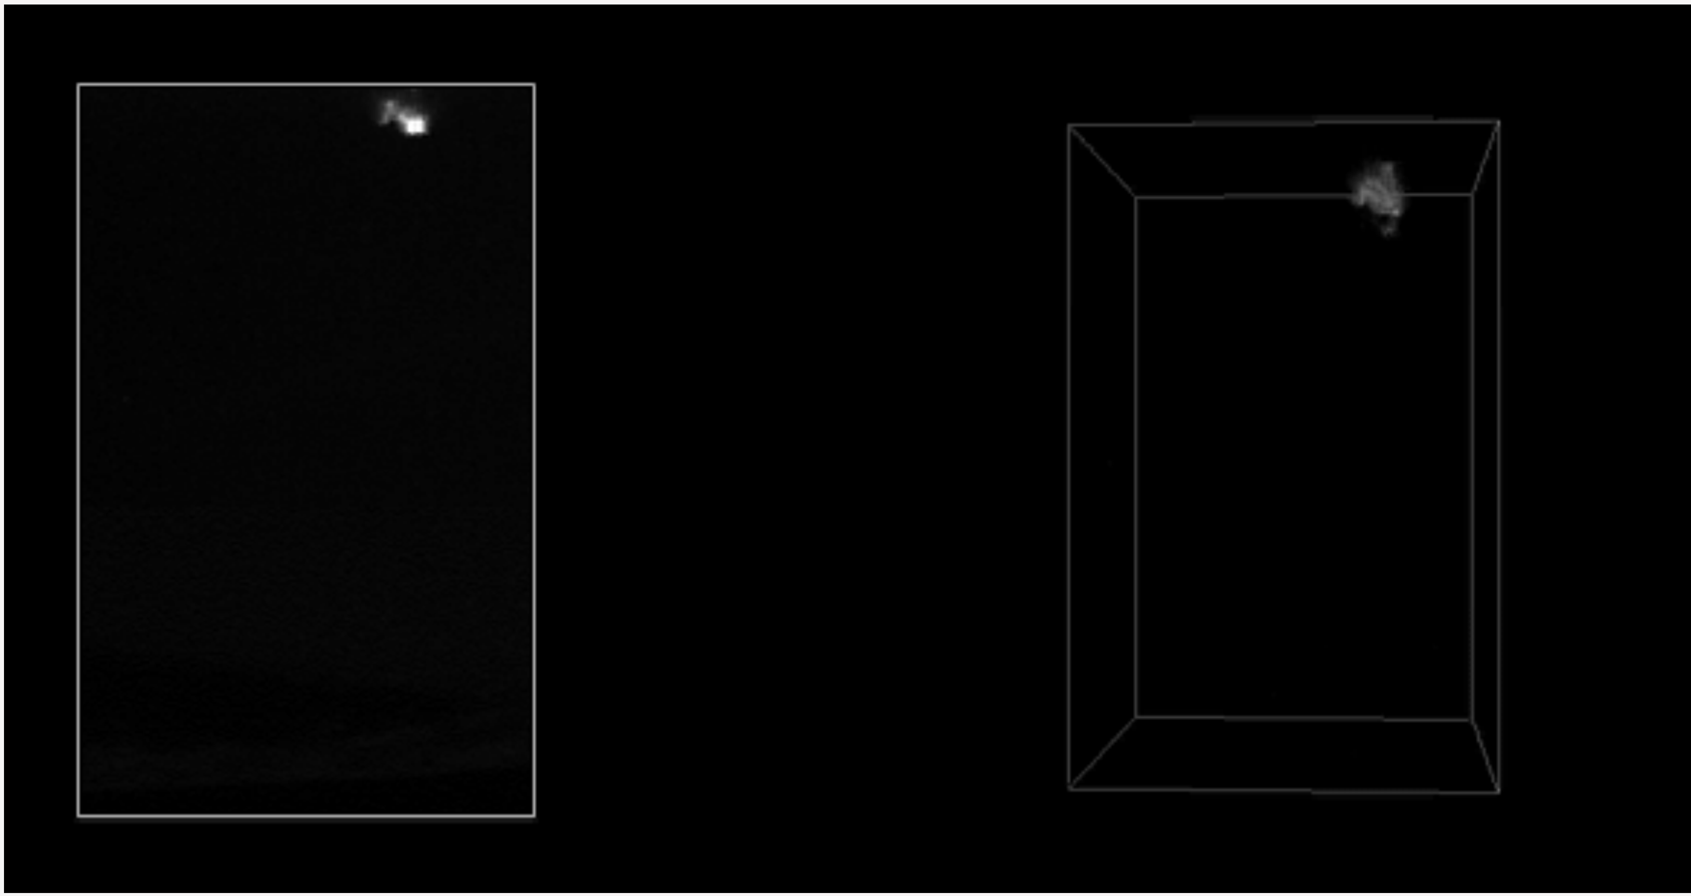

- J. Gu et al, ECCV'08, PAMI'13

Dynamic Volumes – Milk Drop

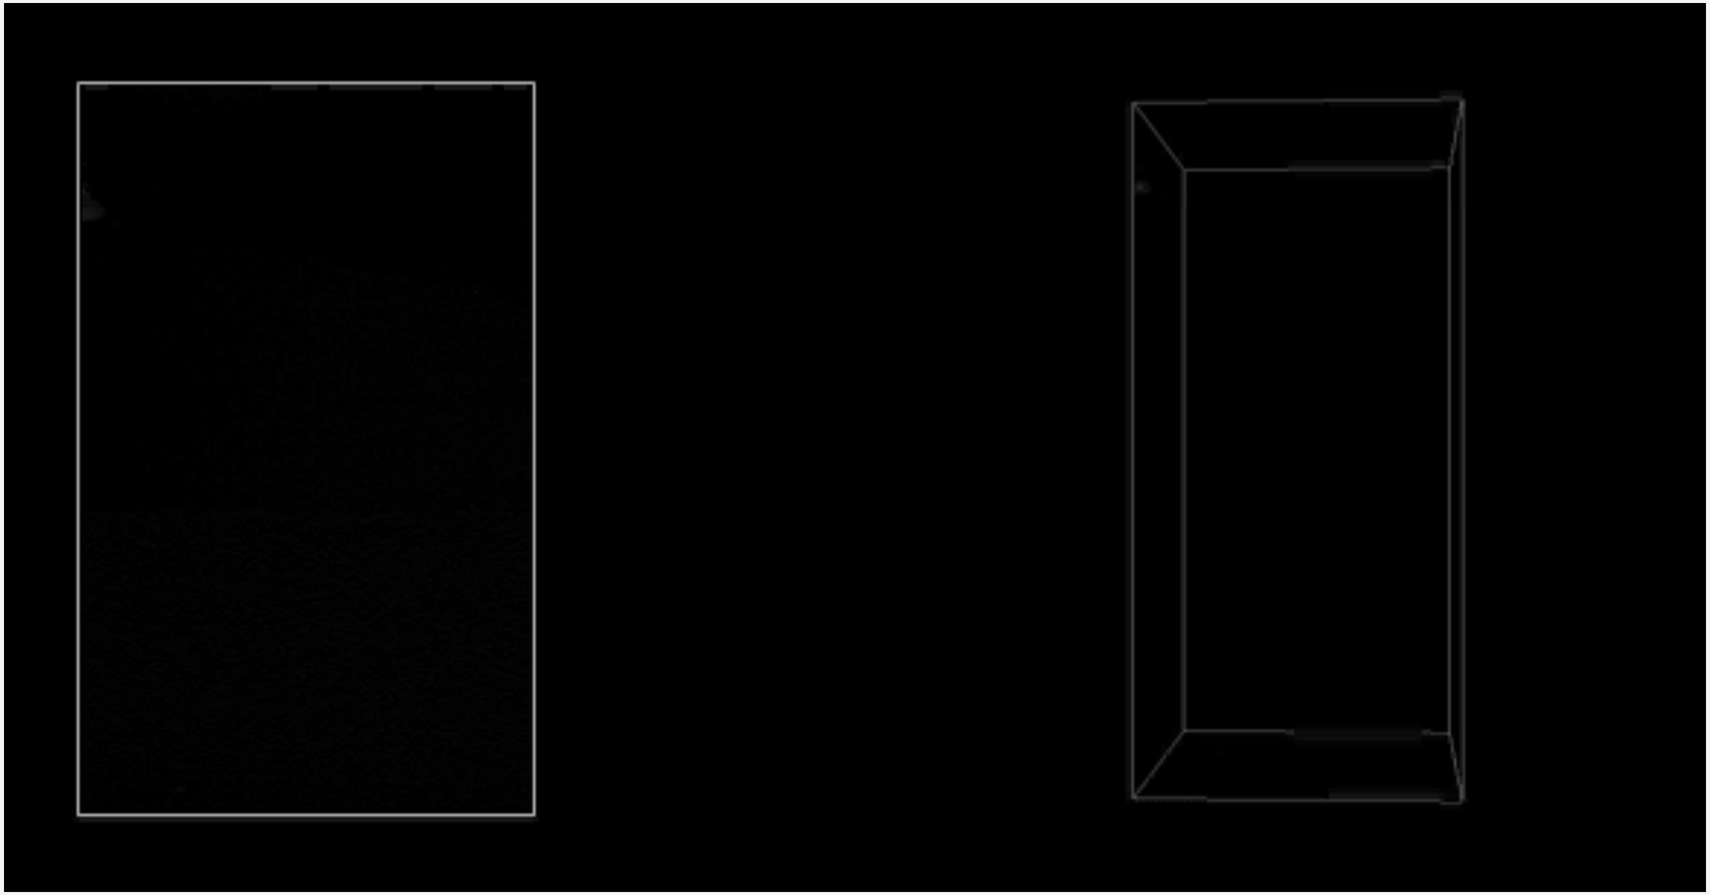

- J. Gu et al, ECCV'08, PAMI'13

Summary

- The property of high sparsity of participating media is exploited through compressive structured light method
  - High efficiency in *acquisition time* and *illumination power*
  - *Attenuation correction* further improves the results

- Limitations
  - *Multiple scattering*: may be alleviated through adopting more complicated light codes
  - Temporal coding is more robust to noise and defocus
  - Spatial coding is preferred due to high temporal resolution, but requires accurate *calibrations*

Conclusion

- Compressive sensing (CS) provides strong theorems for perfect reconstruction of sparse signals
  - For reconstruction of natural images, *informative sensing* is better
- In practice, CS-inspired sparse representation may provide better solutions in many cases
  - However, cannot and will not solve all vision / graphics problems!
- When is possible to apply CS?
  - It is possible to find a *sparse* representation of your data
  - The system has *asymmetrical* computation resources in sensing and reconstruction

Reference

- [1] “Compressed sensing,” IEEE Transactions on Information Theory, 2006
- [2] “Compressive Sensing for Background Subtraction,” ECCV, 2008
- [3] “Informative Sensing of Natural Images,” ICIP 2009
- [4] “Compressive Structured Light for Recovering Inhomogeneous Participating Media,” IEEE Transactions On Pattern Analysis And Machine Intelligence (PAMI), 2013å
- [5] “Robust Face Recognition via Sparse Representation,” IEEE Transactions On Pattern Analysis And Machine Intelligence (PAMI), 2009
- [6] “A Compressive Sensing Approach for Expression-Invariant Face Recognition,” CVPR, 2009
- [7] “Real-time Visual Tracking Using Compressive Sensing,” CVPR, 2011
- [8] “Real-Time Compressive Tracking,” ECCV, 2012
- [9] “CoSaMP: Iterative signal recovery from incomplete and inaccurate samples,” Applied and Computational Harmonic Analysis, 2009