Compressive Sensing for Vision Applications

Richard Baraniuk

Rice University
dsp.rice.edu/cs

Rice University DSP




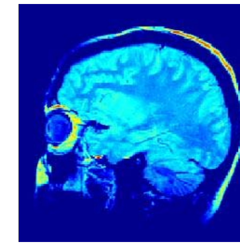
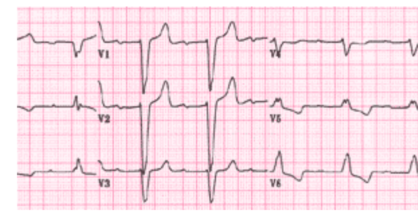
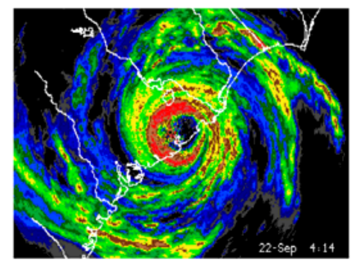

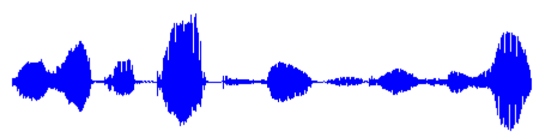


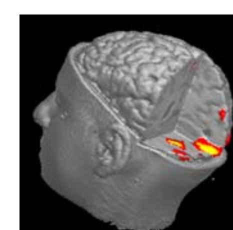

Digital Revolution

Pressure is on DSP

- Success of digital data acquisition is placing increasing pressure on signal/image processing hardware and software to support

higher resolution / denser sampling

» still cameras, video cameras, imaging systems, ...

+

large numbers of sensors

» multi-view image data bases, camera arrays and networks, pattern recognition systems,

+

increasing numbers of modalities

» visual, IR, UV, THz, x-ray, SAR, ...

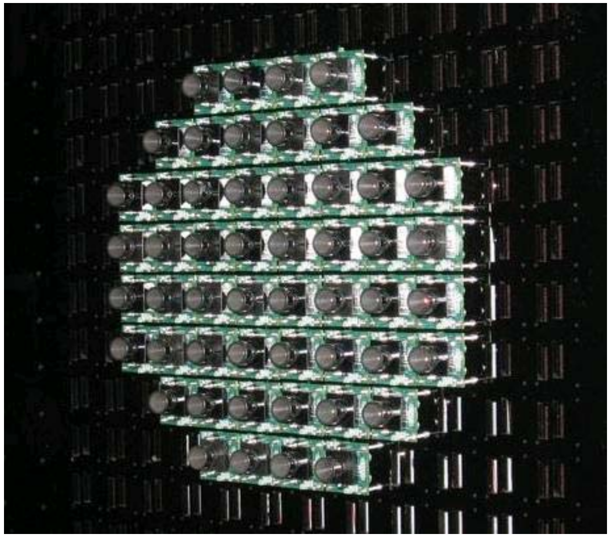

camera arrays

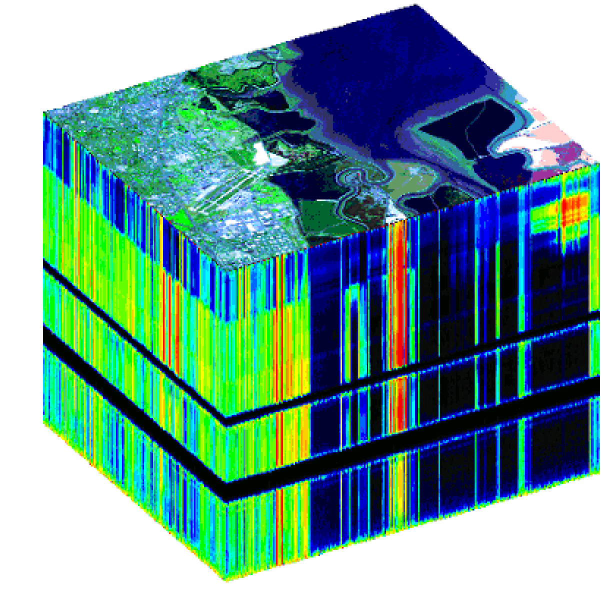

hyperspectral cameras

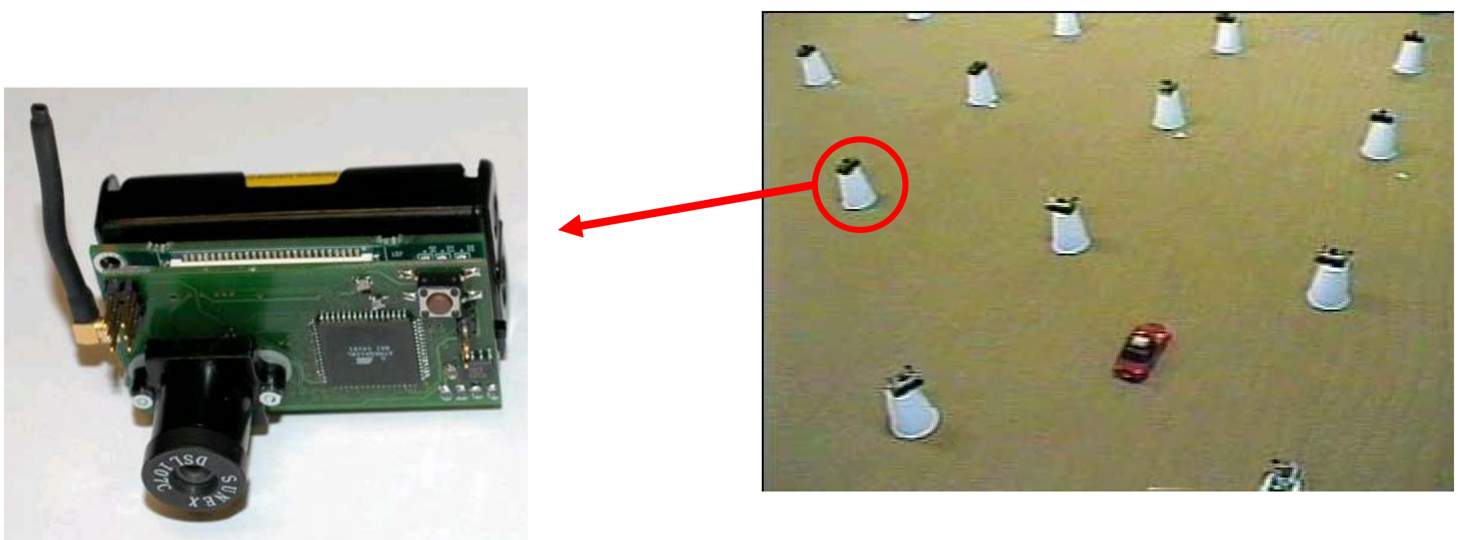

distributed camera networks

Pressure is on DSP

- Success of digital data acquisition is placing increasing pressure on signal/image processing hardware and software to support

higher resolution / denser sampling

» still cameras, video cameras, imaging systems, ...

+

large numbers of sensors

» multi-view image data bases, camera arrays and networks, pattern recognition systems,

+

increasing numbers of modalities

» visual, IR, UV, THz, x-ray, ...

=

deluge of data

» how to acquire, store, fuse, process efficiently?

Antipasto

Sensing by *Sampling*

Sensing by *Sampling*

- Long-established paradigm for digital data acquisition
  - ***sample*** data at Nyquist rate (2x bandwidth)
  - ***compress*** data (signal-dependent, nonlinear)

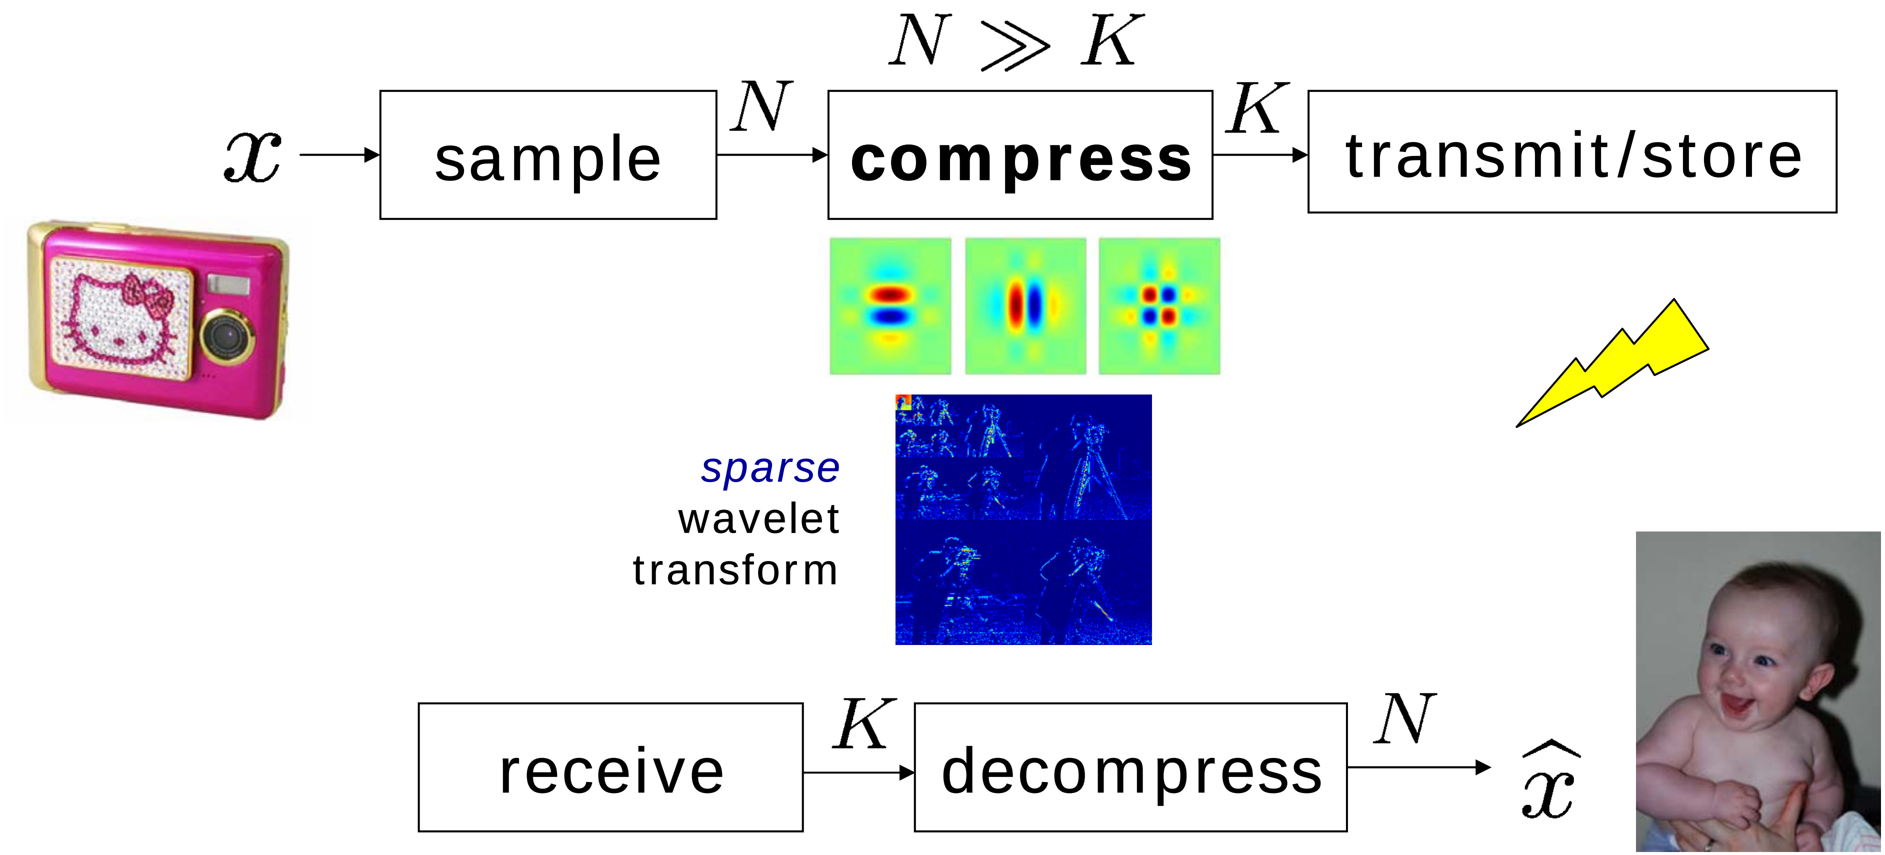

Sparsity / Compressibility

N
pixels

$K \ll N$
large
wavelet
coefficients

N
wideband
signal
samples

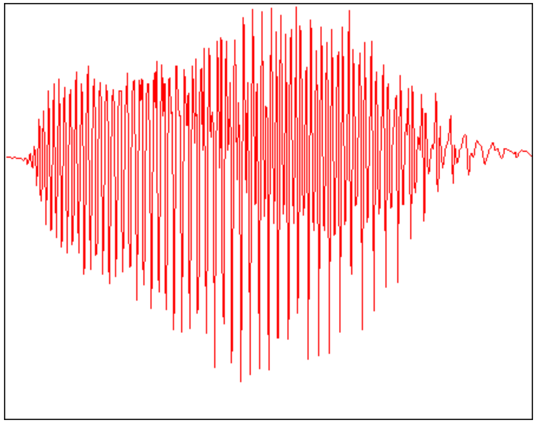

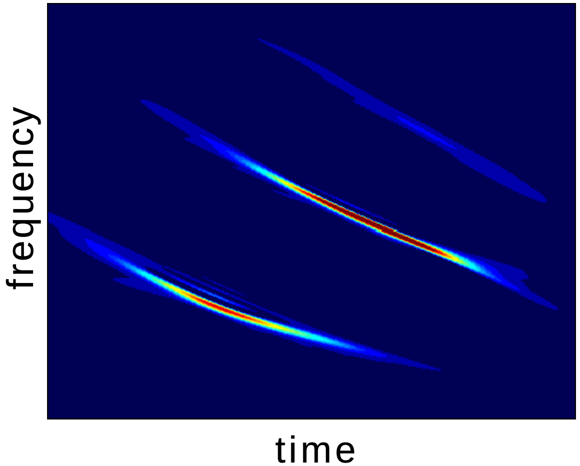

$K \ll N$
large
Gabor
coefficients

What's Wrong with this Picture?

- Long-established paradigm for digital data acquisition
  - ***sample*** data at Nyquist rate (2x bandwidth)
  - ***compress*** data (signal-dependent, nonlinear)
  - ***brick wall*** to resolution/performance

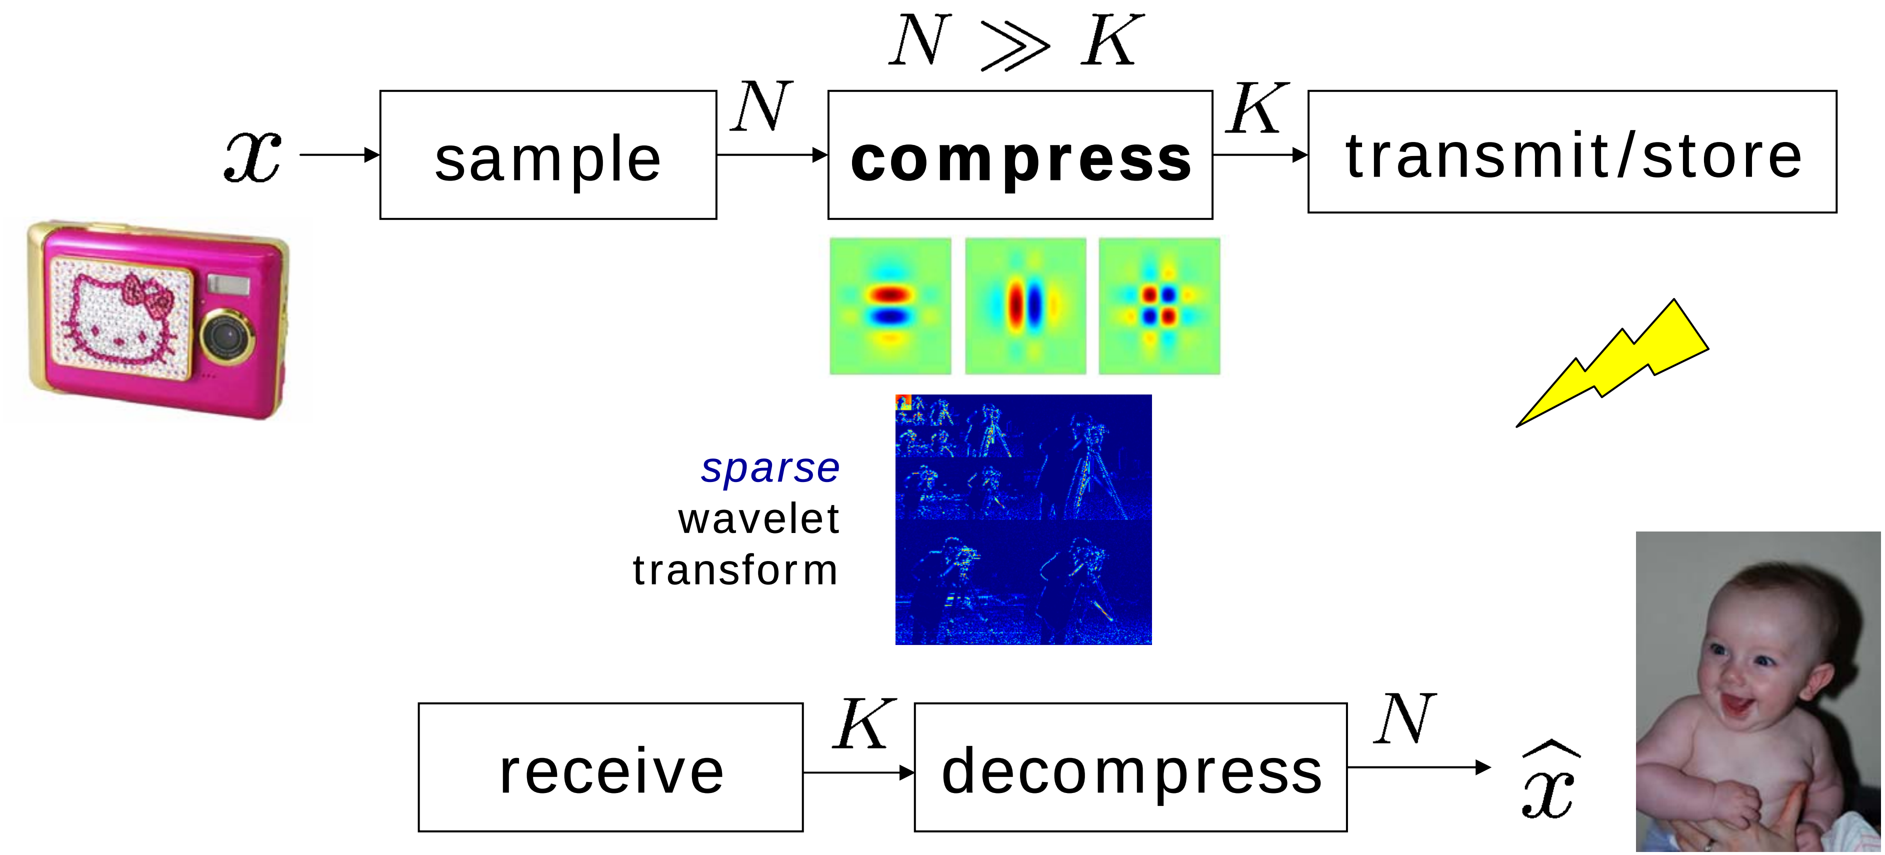

Primo

Compressive Sensing

Compressive Sensing (CS)

- Recall Shannon/Nyquist theorem
  - Shannon was a *pessimist*
  - 2x oversampling Nyquist rate is a worst-case bound for *any* bandlimited data
  - sparsity/compressibility irrelevant
  - Shannon sampling is a linear process while compression is a nonlinear process

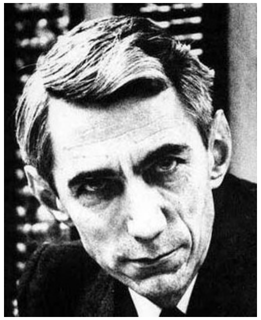

- **Compressive sensing**
  - new sampling theory that *leverages compressibility*
  - based on new *uncertainty principles*
  - *randomness* plays a key role

Compressive Sensing

- Directly acquire "*compressed*" data
- Replace samples by more general "measurements"

$$K < \underline{M} \ll N$$

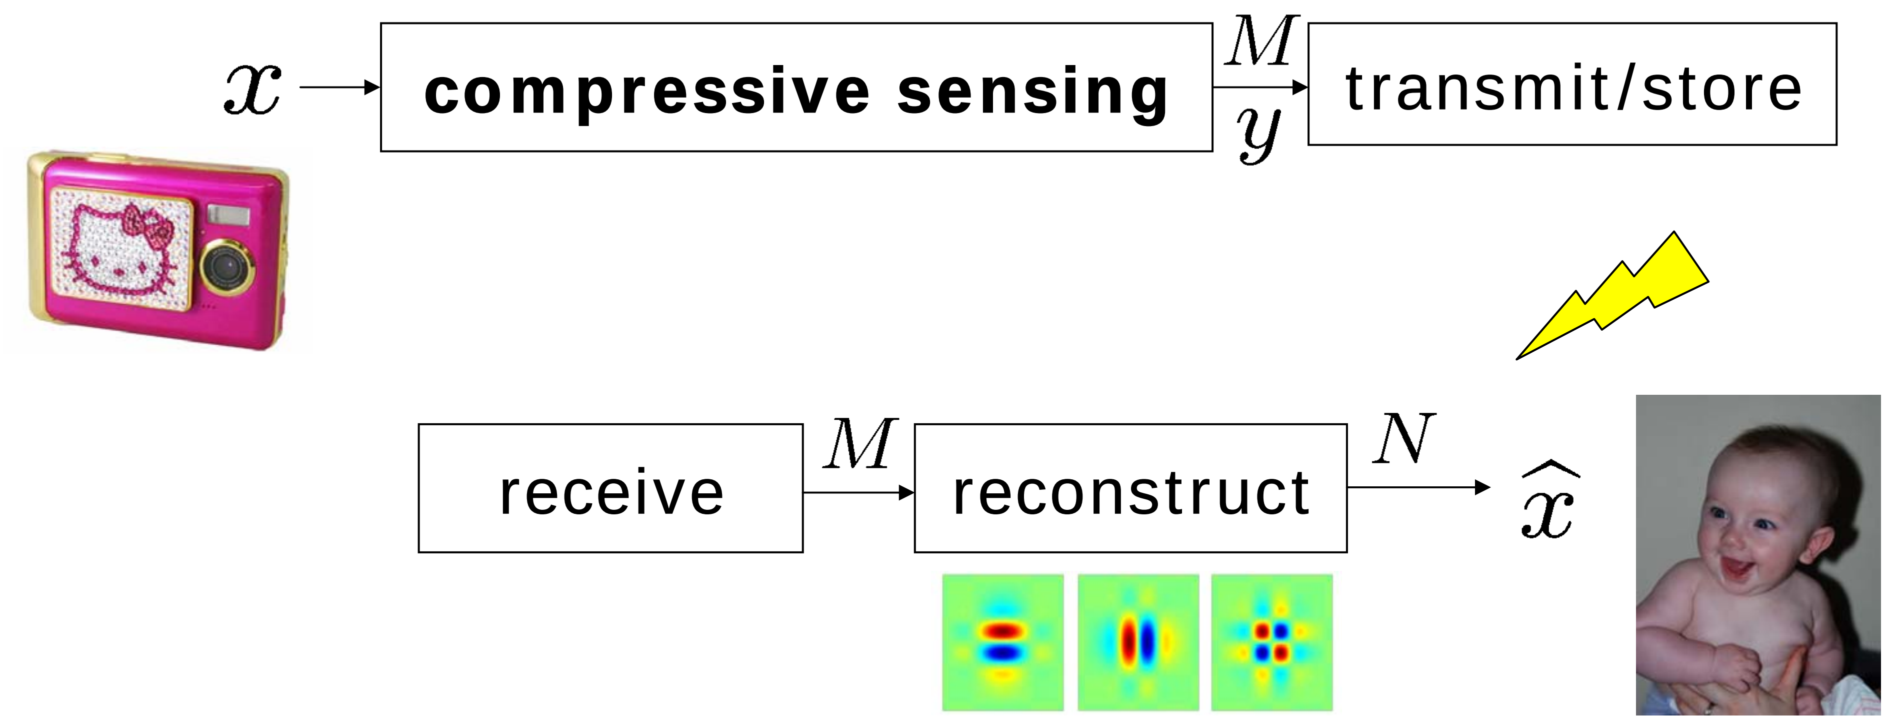

Sampling

- Signal x is K-*sparse* in basis/dictionary Ψ
  - WLOG assume sparse in space domain $\Psi = I$
- **Samples**

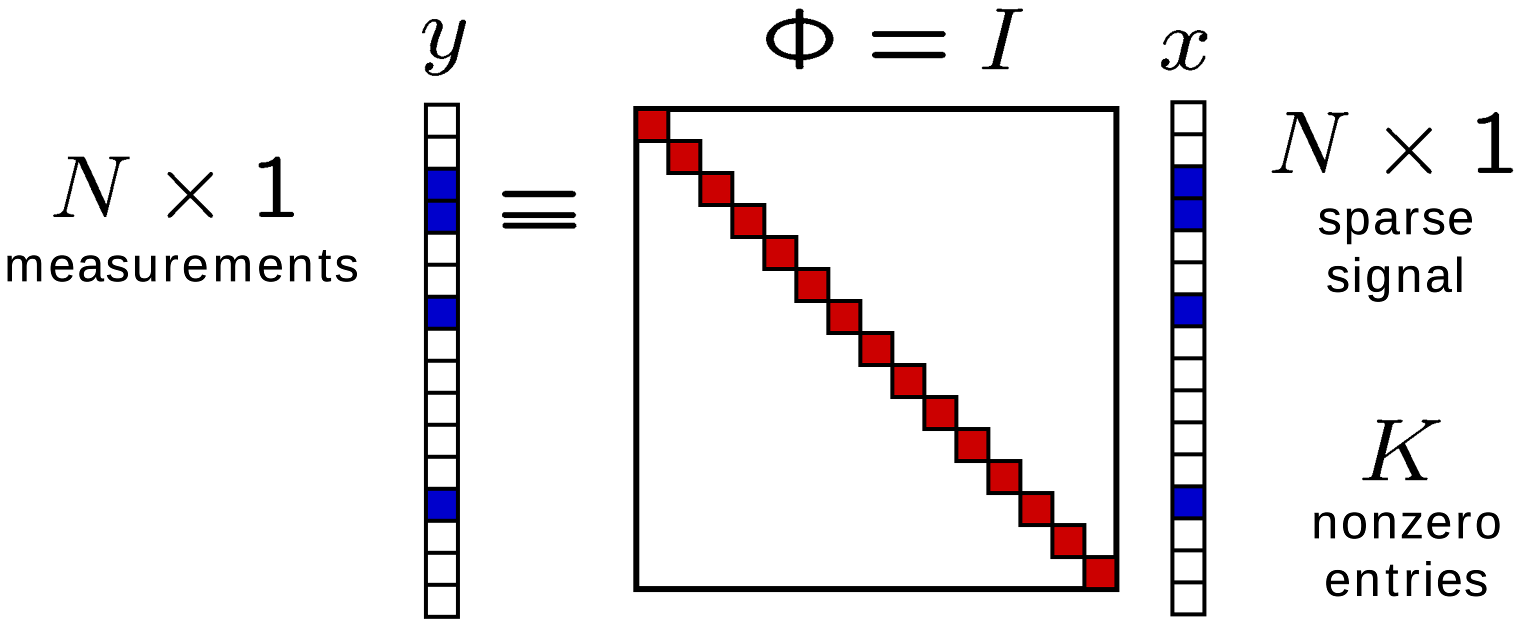

Compressive Data Acquisition

- When data is sparse/compressible, can directly acquire a ***condensed representation*** with no/little information loss through ***dimensionality reduction***

$$y = \Phi x$$

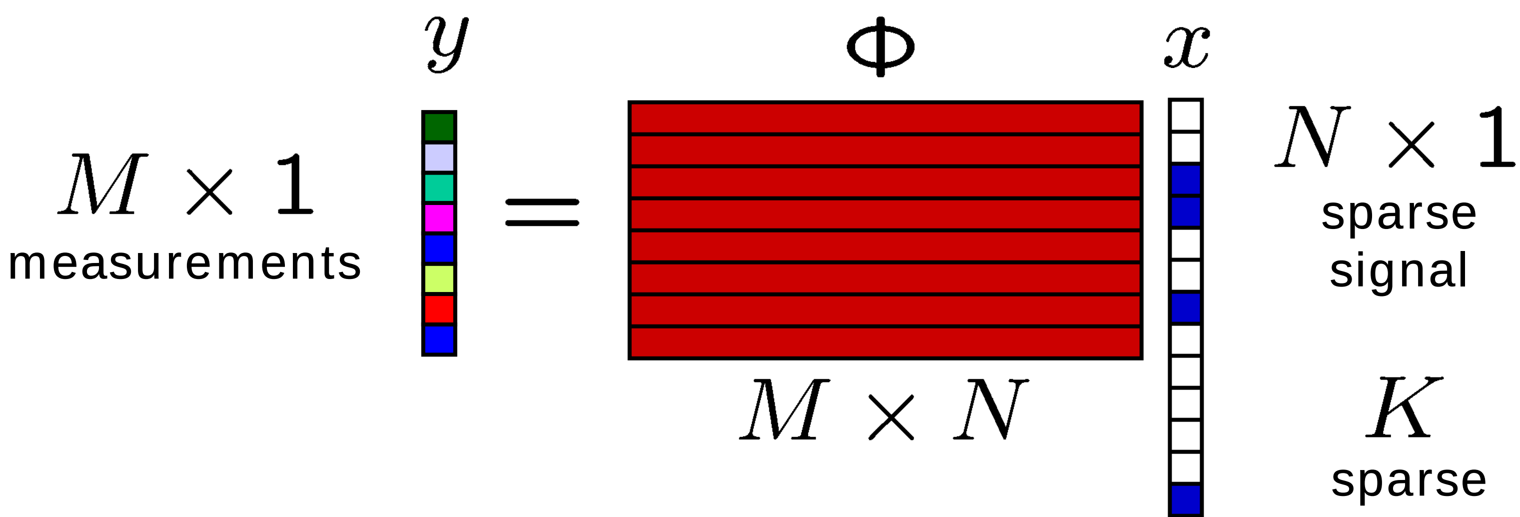

$$K < M \ll N$$

Compressive Data Acquisition

- When data is sparse/compressible, can directly acquire a ***condensed representation*** with no/little information loss

$$y = \Phi x$$

- ***Random projection*** will work

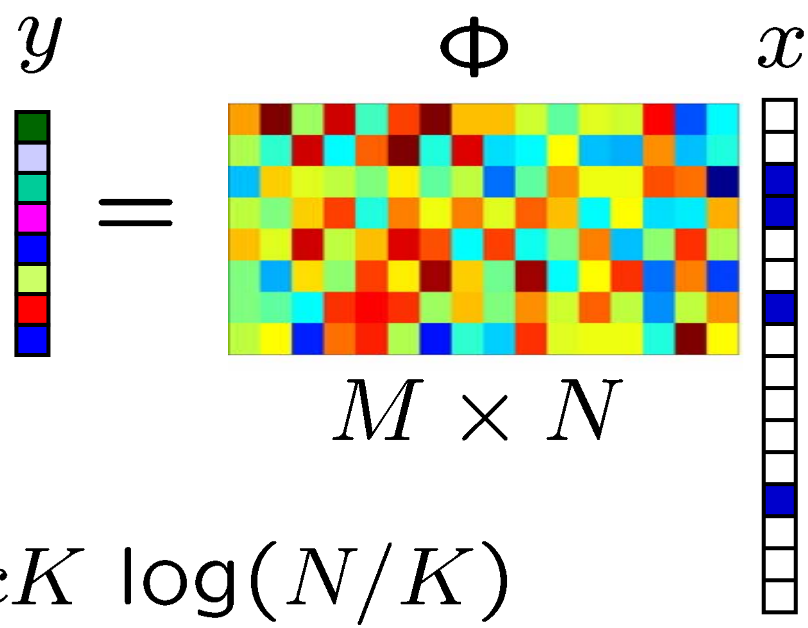

$$M = cK \log(N/K)$$

[Candes-Romberg-Tao, Donoho, 2004]

Compressive Sensing

- Directly acquire “***compressed***” data
- Replace samples by more general “measurements”

$$M = cK \log(N/K)$$

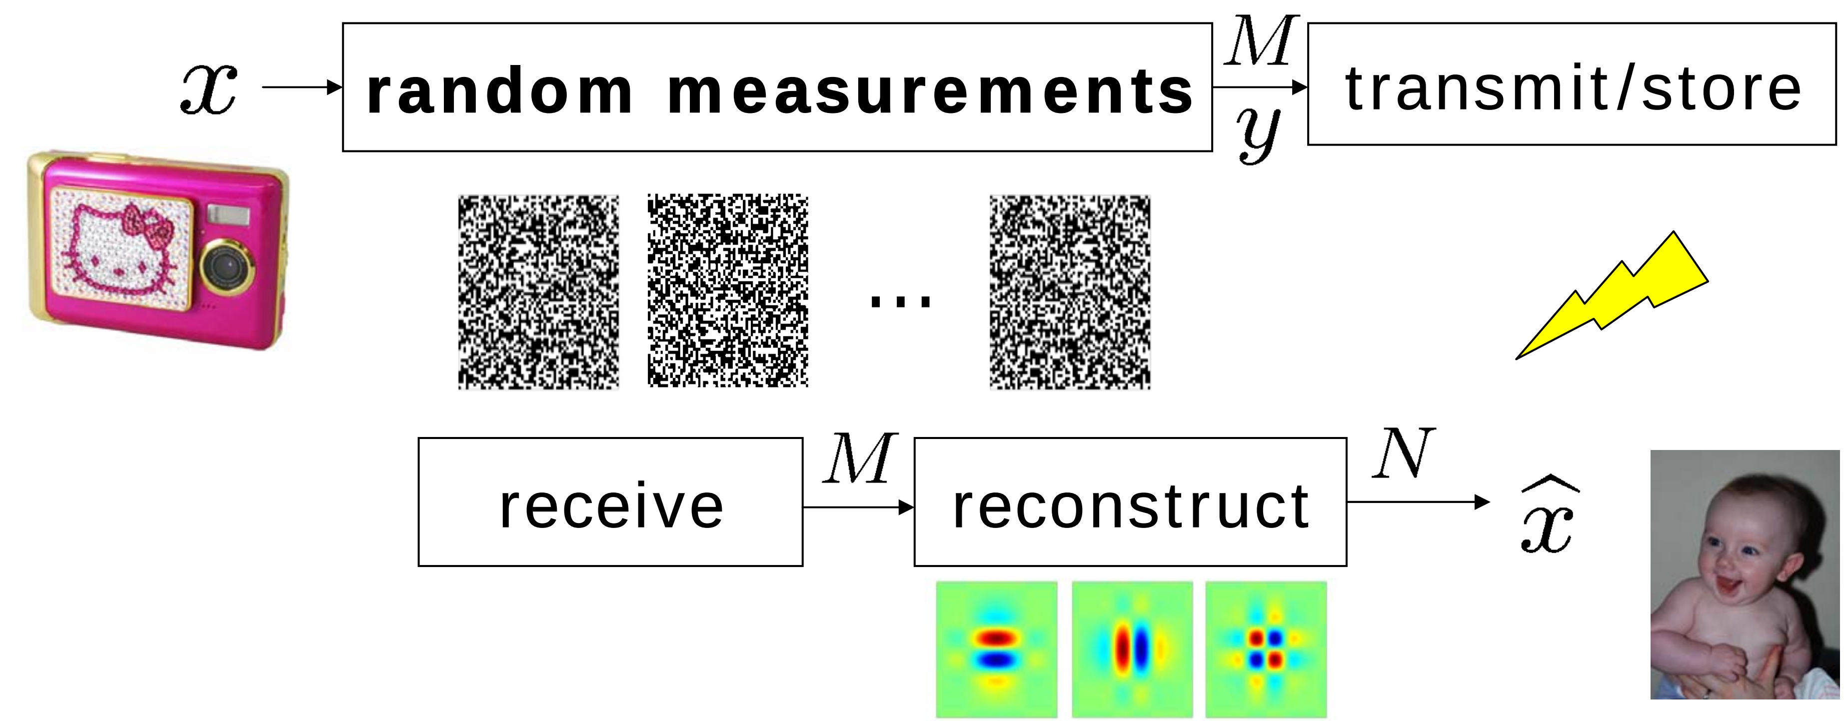

CS Signal Recovery

- Reconstruction/decoding: given $y = \Phi x$
  (**ill-posed inverse problem**) find x

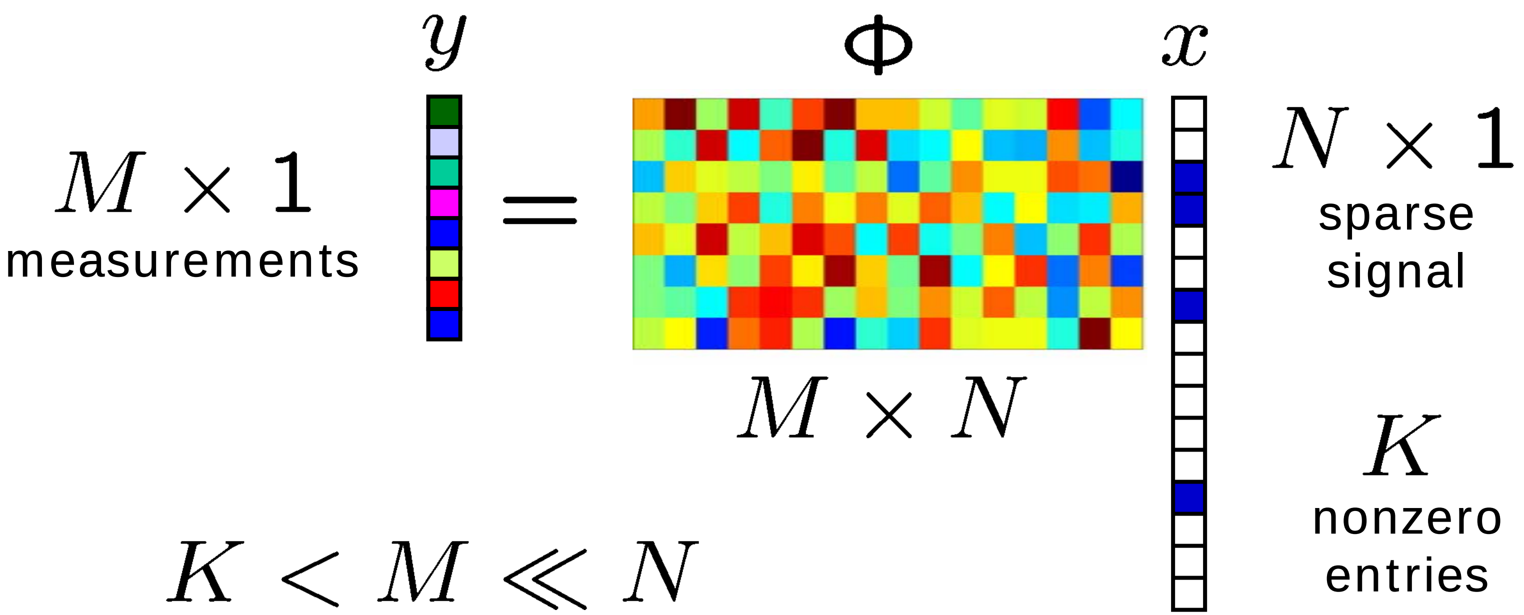

$$K < M \ll N$$

CS Signal Recovery

- Reconstruction/decoding: given $y = \Phi x$
  (ill-posed inverse problem) find x

- **Null space**: there are infinitely many x
  such that $y = \Phi x$

- So search in null space for the "best" x according to some criterion
  - ex: least squares

CS Signal Recovery

- Reconstruction/decoding: given $y = \Phi x$
  (ill-posed inverse problem) find x

- ***L*₂** fast, **wrong** $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_2$

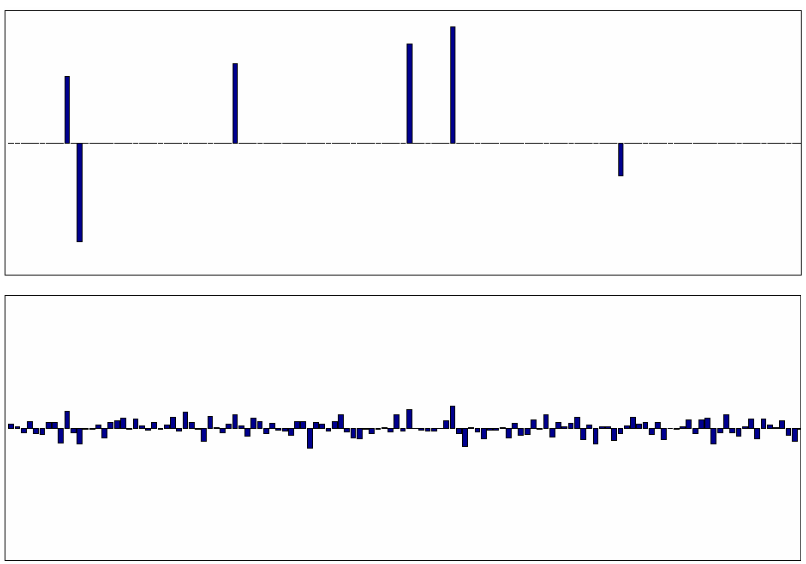

x

$\widehat{x} = (\Phi^T \Phi)^{-1} \Phi^T y$

CS Signal Recovery

- Reconstruction/decoding: given $y = \Phi x$
  (ill-posed inverse problem) find x

- ***L_2*** fast, wrong $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_2$

- ***L_0*** **correct, slow**
  only M=K+1 measurements required to perfectly reconstruct K-sparse signal
  [Bresler; Rice]

$$\widehat{x} = \arg\min_{y=\Phi x} \|x\|_0$$

number of nonzero entries

CS Signal Recovery

- Reconstruction/decoding: given $y = \Phi x$
  (ill-posed inverse problem) find x

- ***L*$_2$** fast, wrong $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_2$

- ***L*$_0$** correct, slow $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_0$

- ***L*$_1$** **correct,**
  mild oversampling $\widehat{x} = \arg\min_{y=\Phi x} \|x\|_1$
  [Candes, Romberg, Tao; Donoho] ***linear program***

$$M = cK \log(N/K) \ll N$$



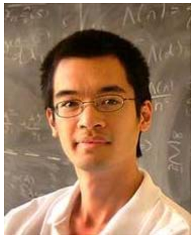


Compressive Sensing

- Directly acquire "***compressed***" data
- Replace samples by more general "measurements"

$$M = cK \log(N/K)$$

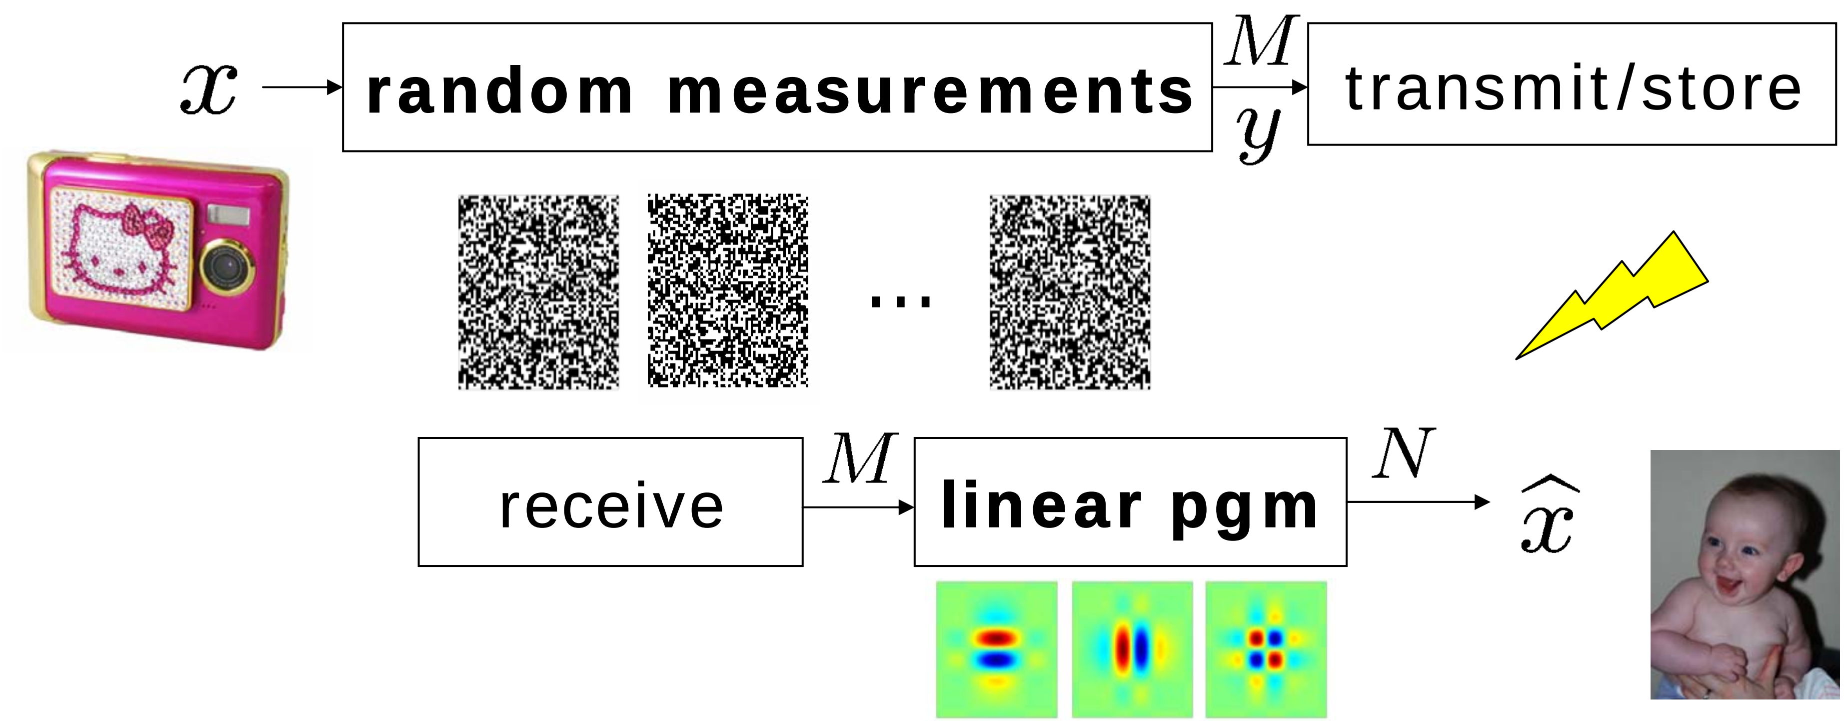

CS Signal Recovery

original (65k pixels)

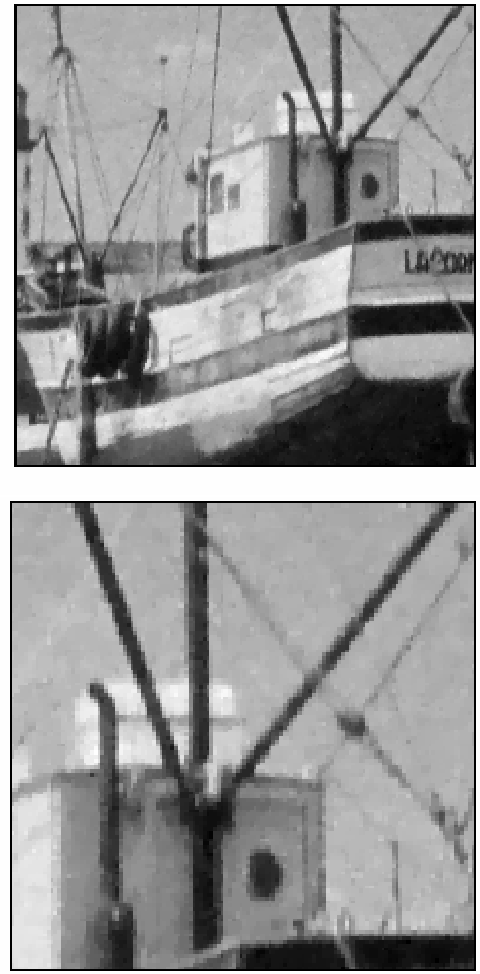

20k random projections

7k–term wavelet approximation

CS Hallmarks

- CS changes the rules of the data acquisition game
  - exploits a priori signal *sparsity* information
- **Universal**
  - same random projections / hardware can be used for *any* compressible signal class (*generic*)
- **Democratic**
  - each measurement carries the same amount of information
  - simple encoding
  - robust to measurement loss and quantization
- **Asymmetrical** (most processing at decoder)
- Random projections weakly **encrypted**

Secondo

Compressive Sensing
in Action

“Single-Pixel” CS Camera

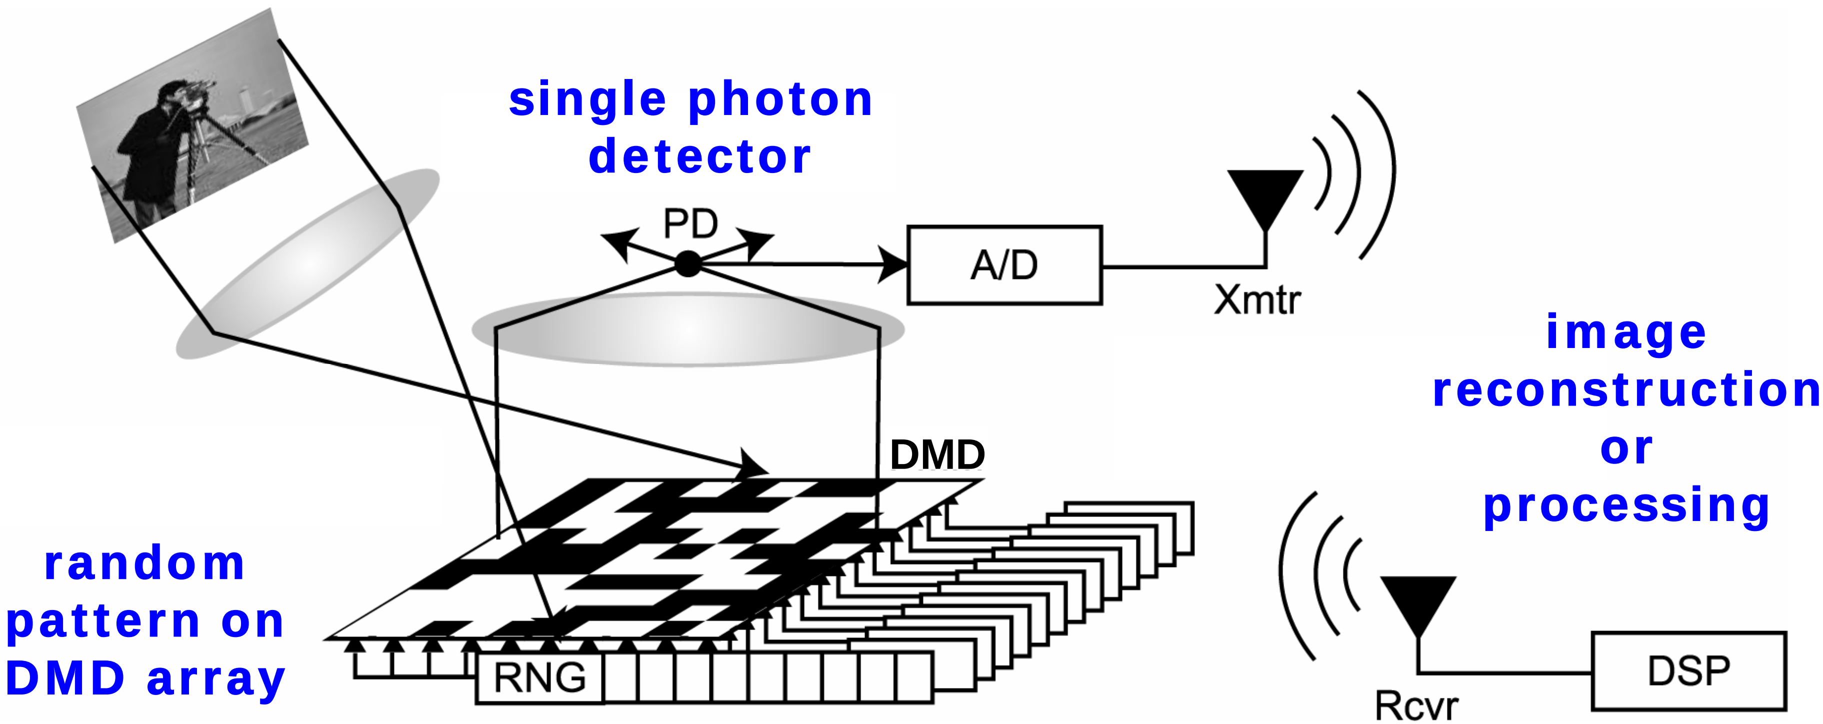

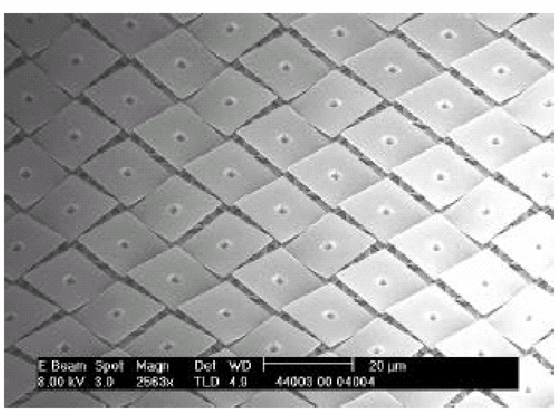

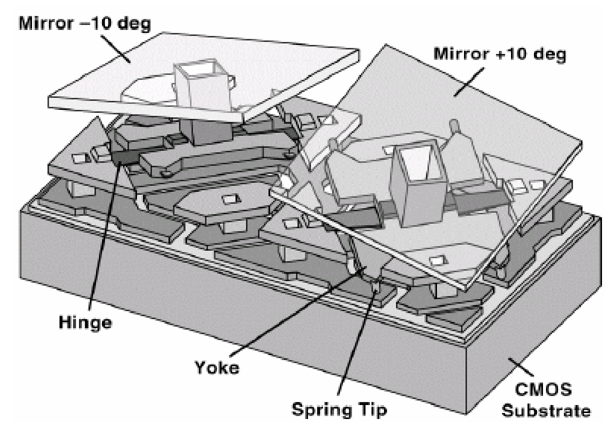

w/ Kevin Kelly
and students

“Single-Pixel” CS Camera

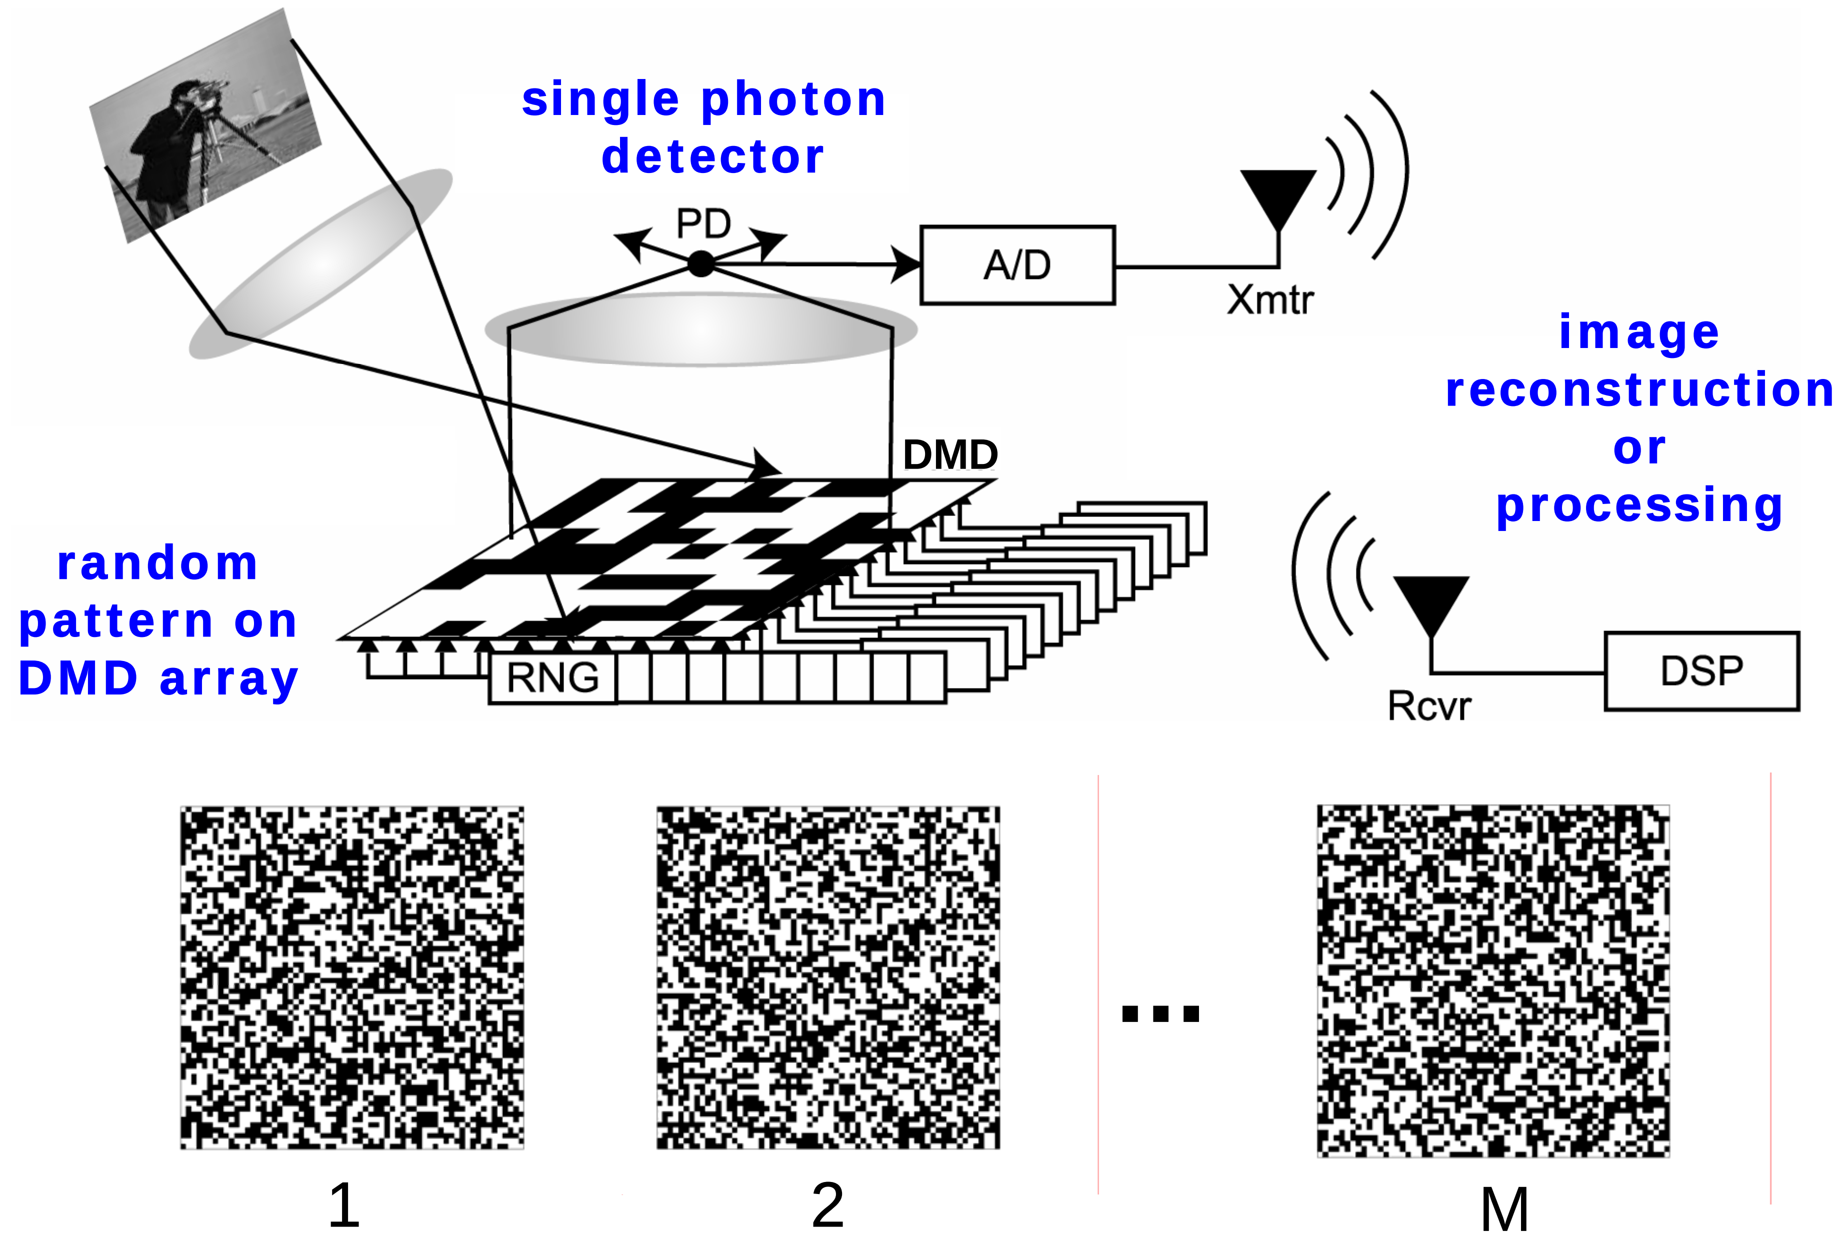

1 Chip DLP™ Projection

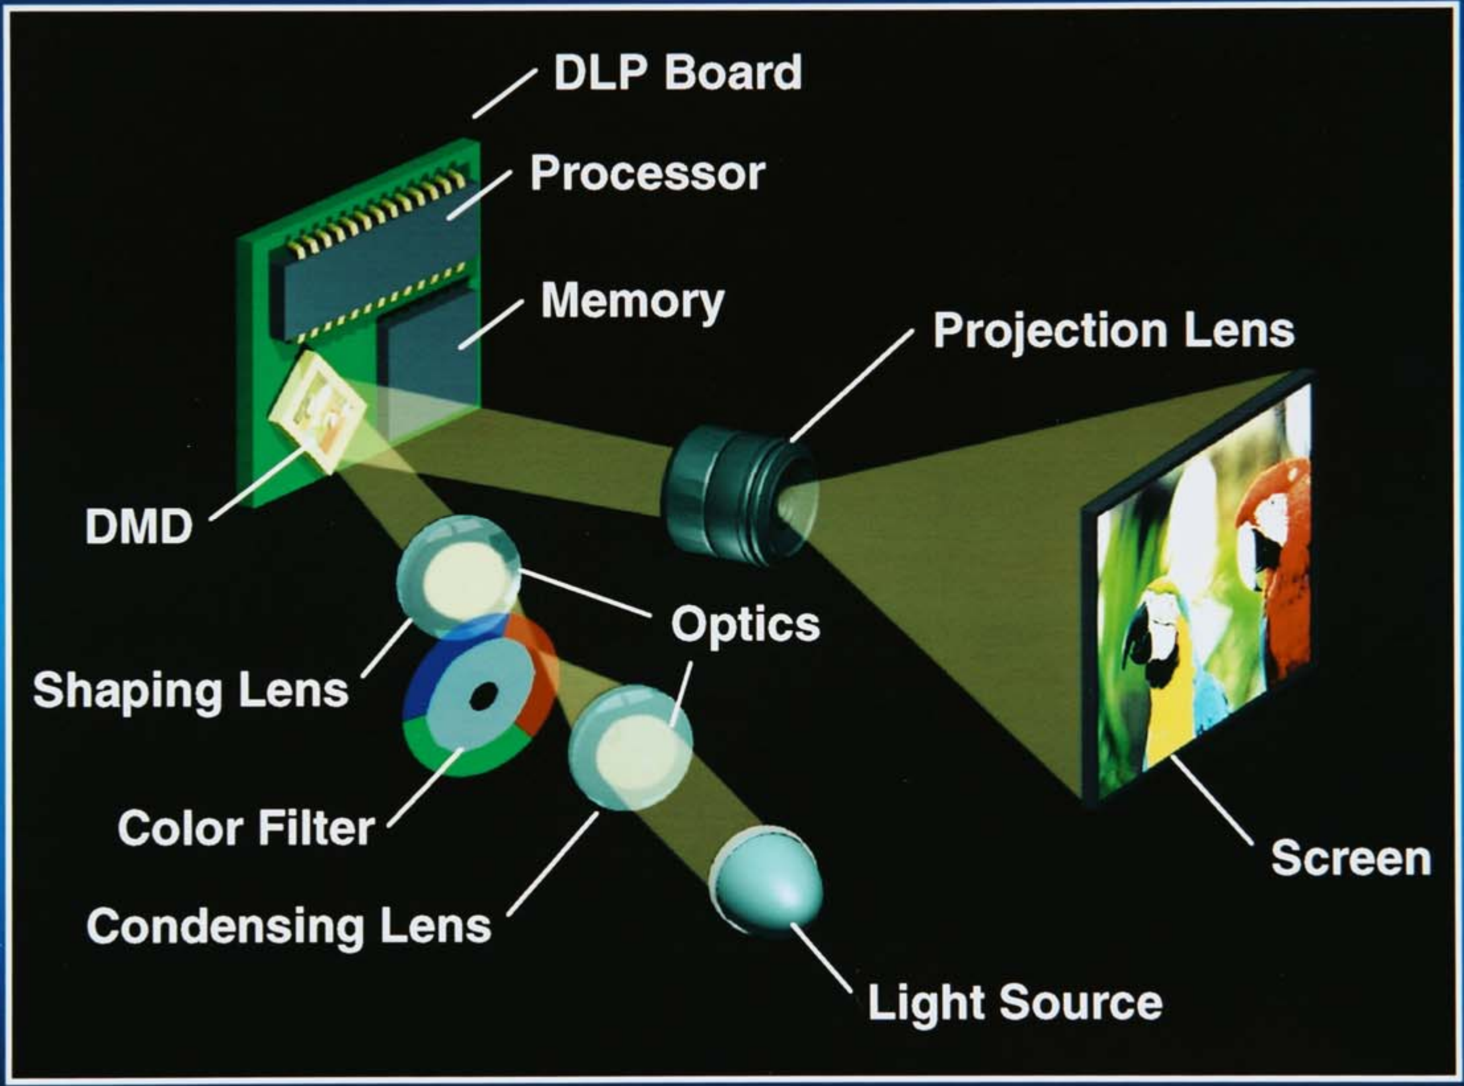

Single Pixel Camera

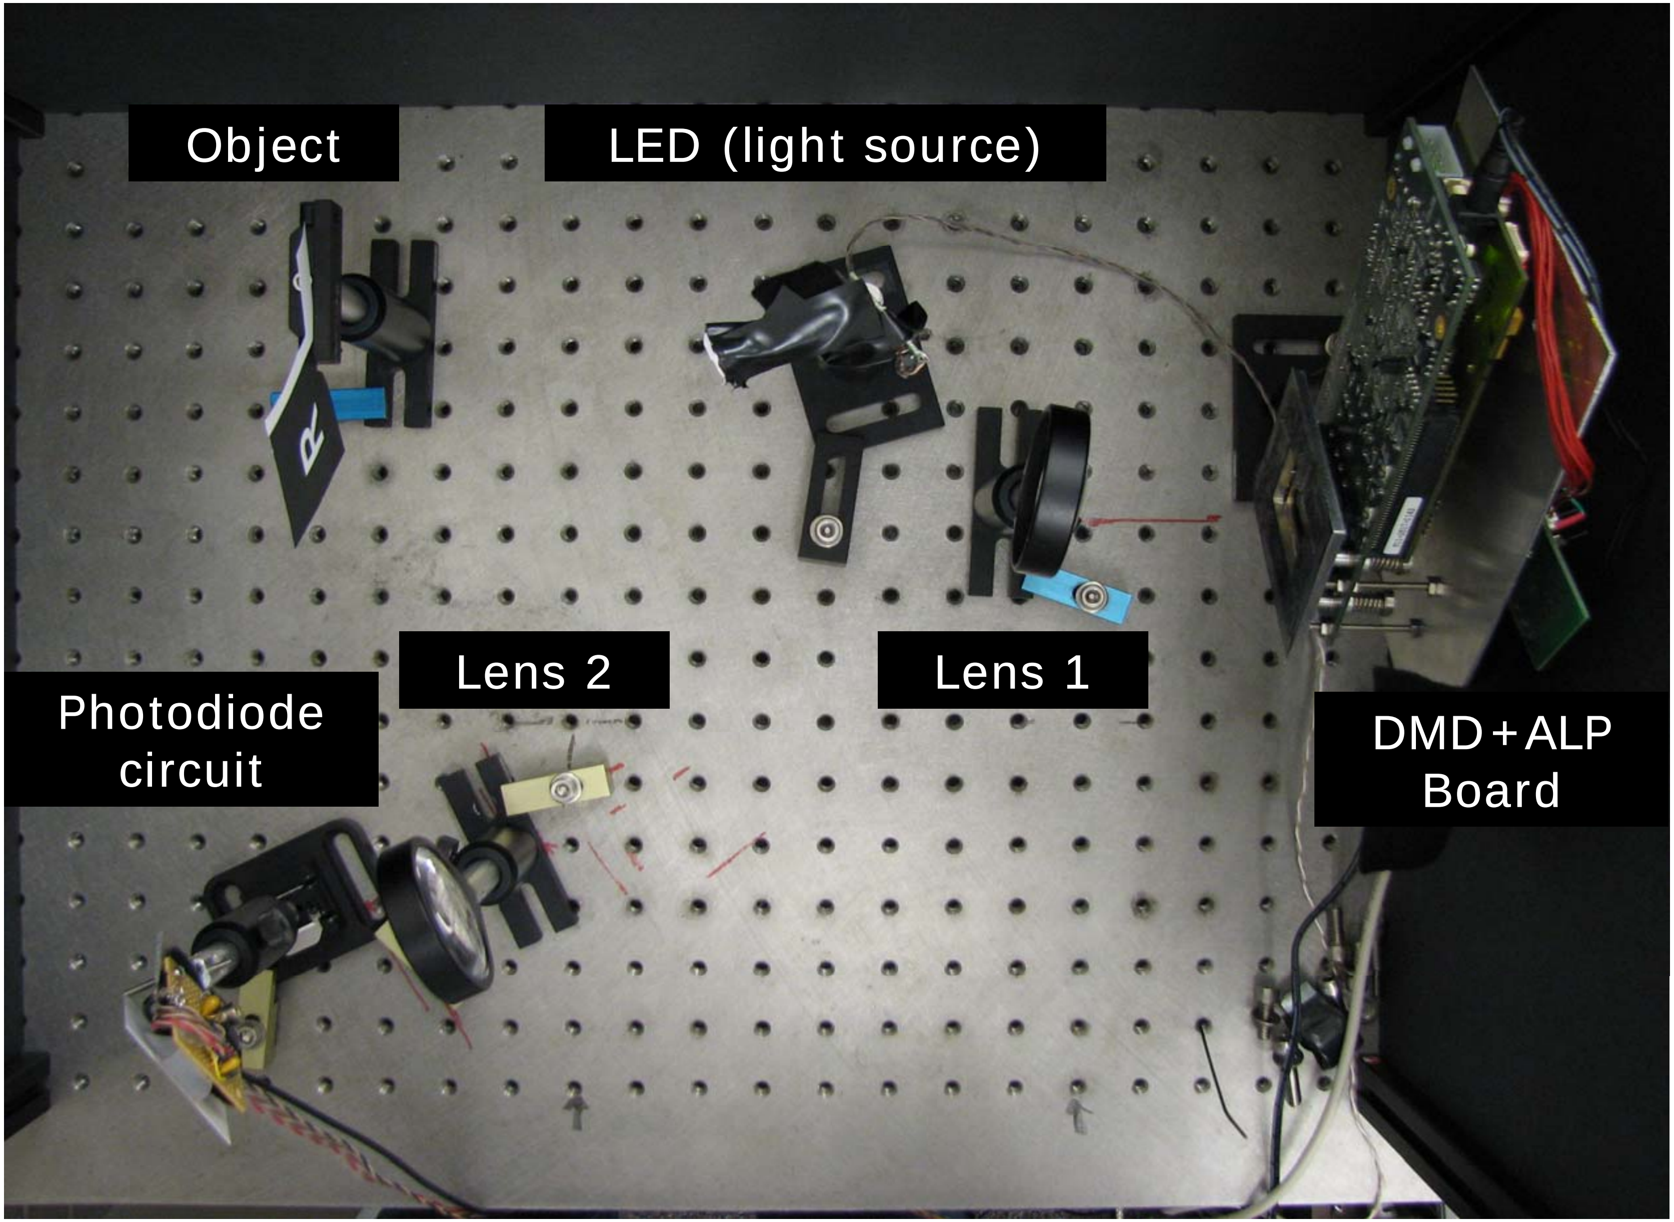

First Image Acquisition

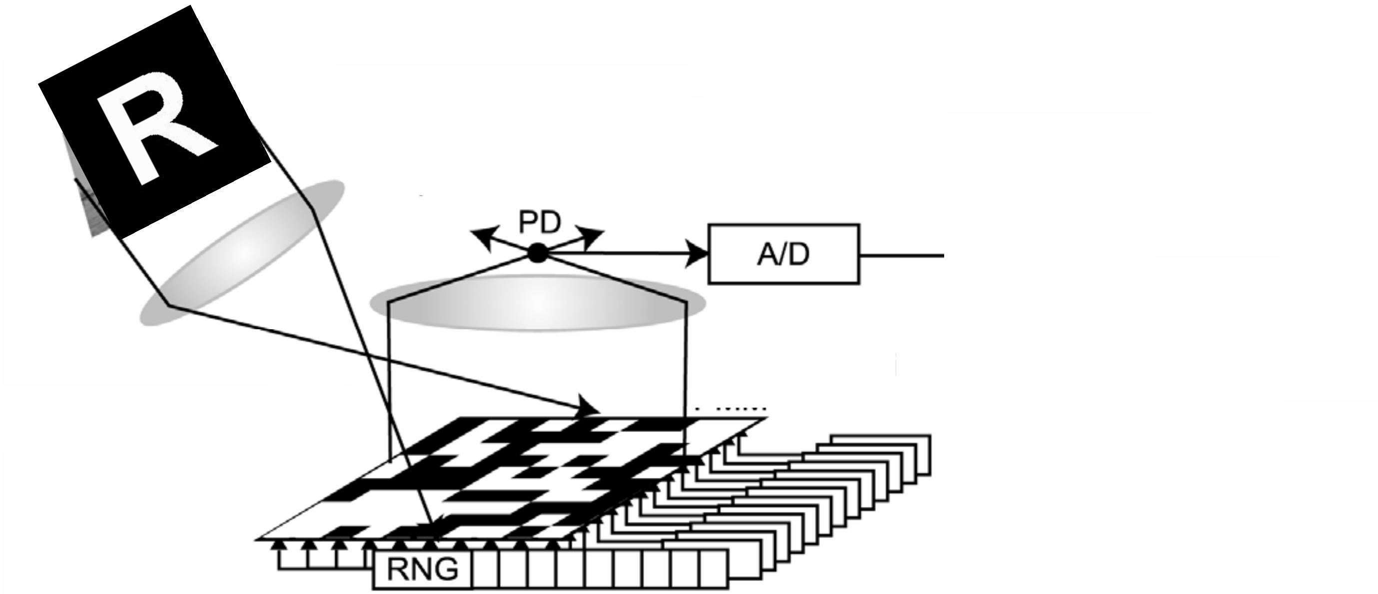

target
65536 pixels

11000 measurements
(16%)

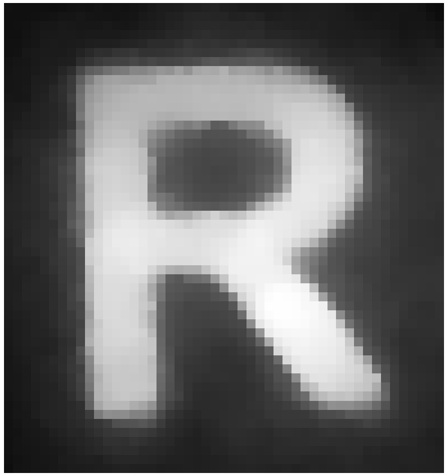

1300 measurements
(2%)

World's First Photograph

- 1826, Joseph Niepce
- Farm buildings and sky
- 8 hour exposure
- On display at UT-Austin

Second Image Acquisition

4096
pixels

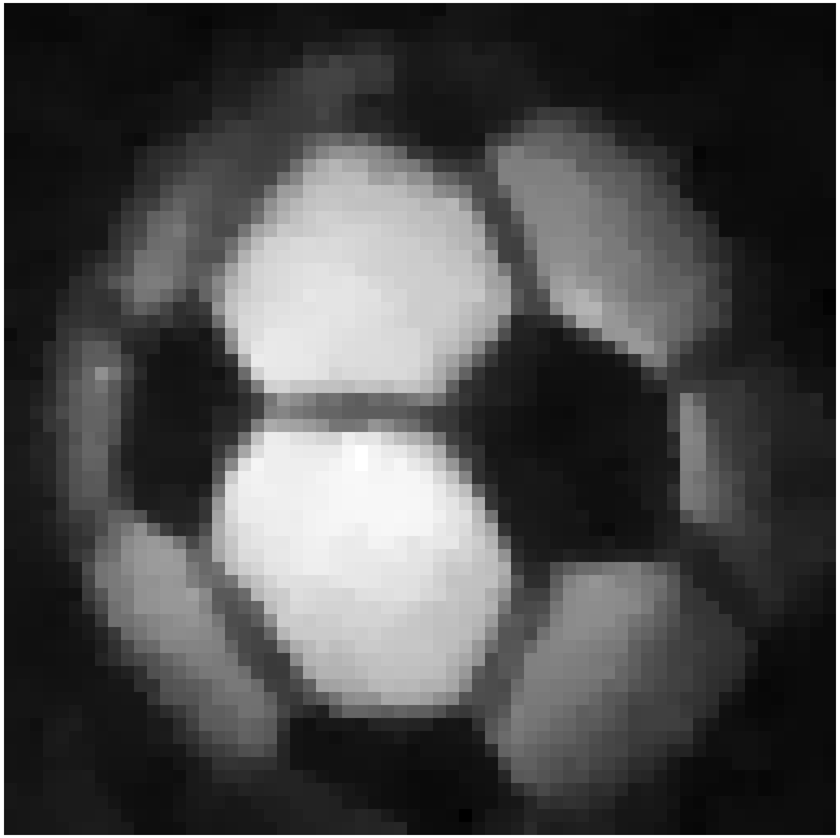

500
random measurements

Single-Pixel Camera

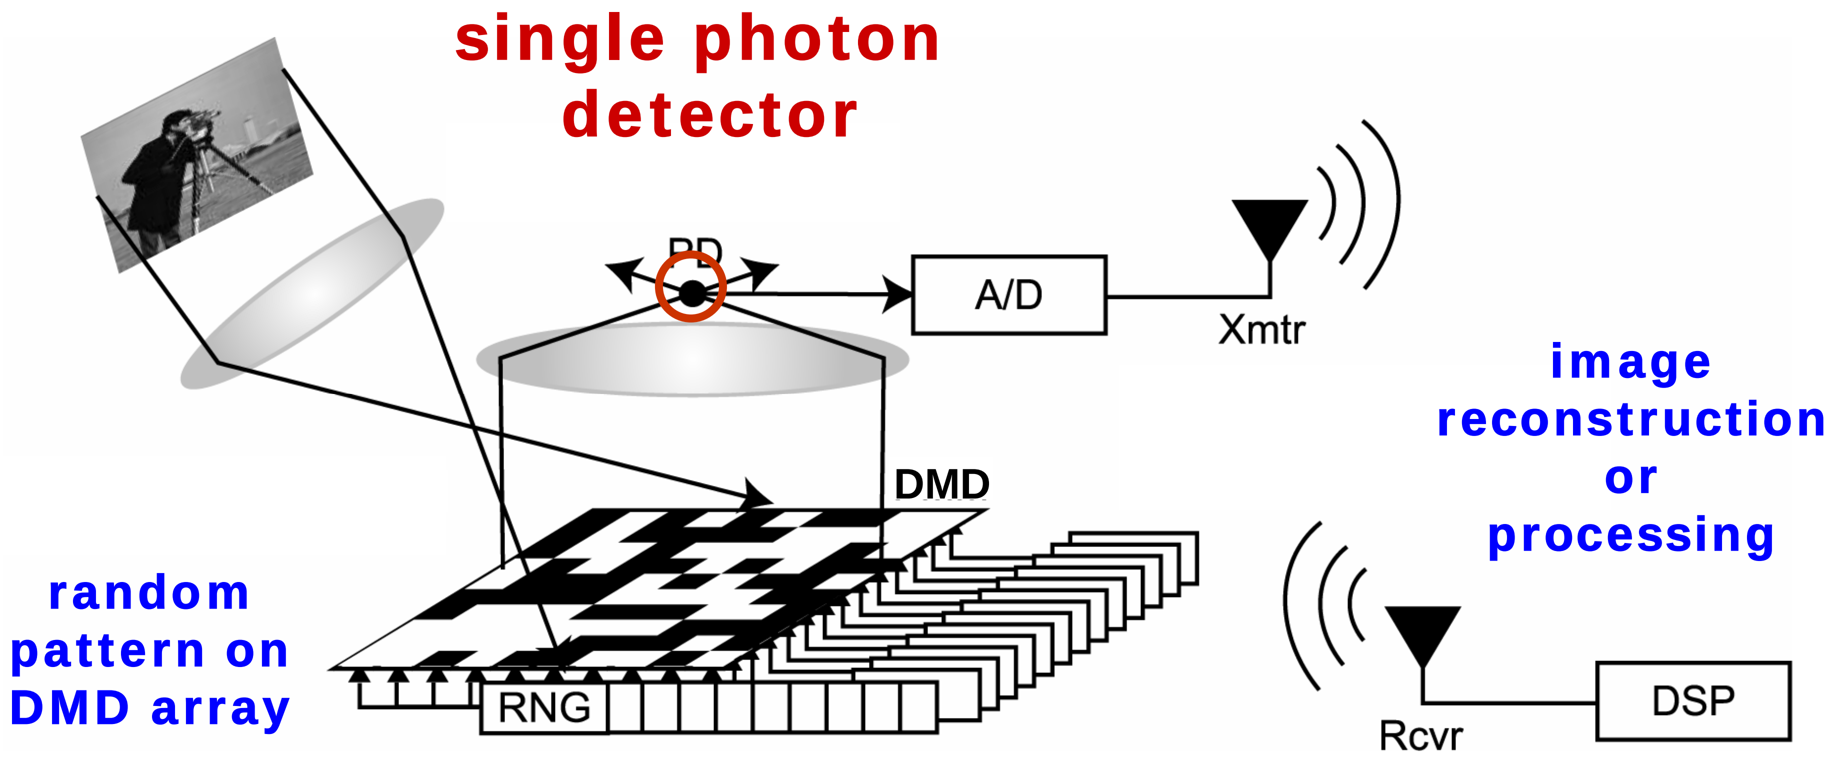

Dual Visible and Infrared Imaging

ADVANCED PHOTONIX, INC.®

SD138-11-31-211
Silicon PIN Photodiode Sandwich Detector

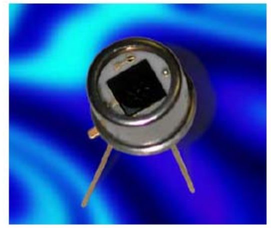

dual photodiode sandwich

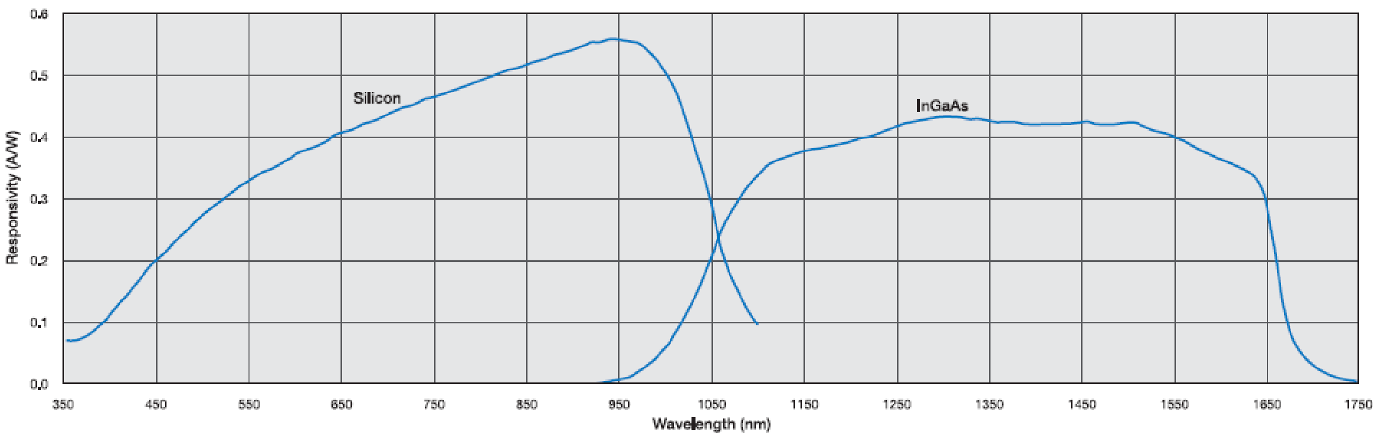

K cutout in paper

front-lit
visible

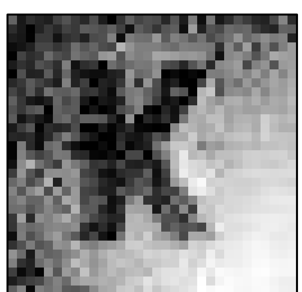

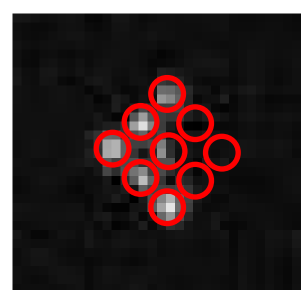

back-lit
IR LEDs

Dual Visible and Infrared Imaging

ADVANCED PHOTONIX, INC.®

SD138-11-31-211
Silicon PIN Photodiode Sandwich Detector

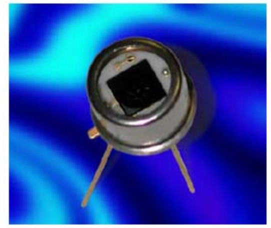

dual photodiode sandwich

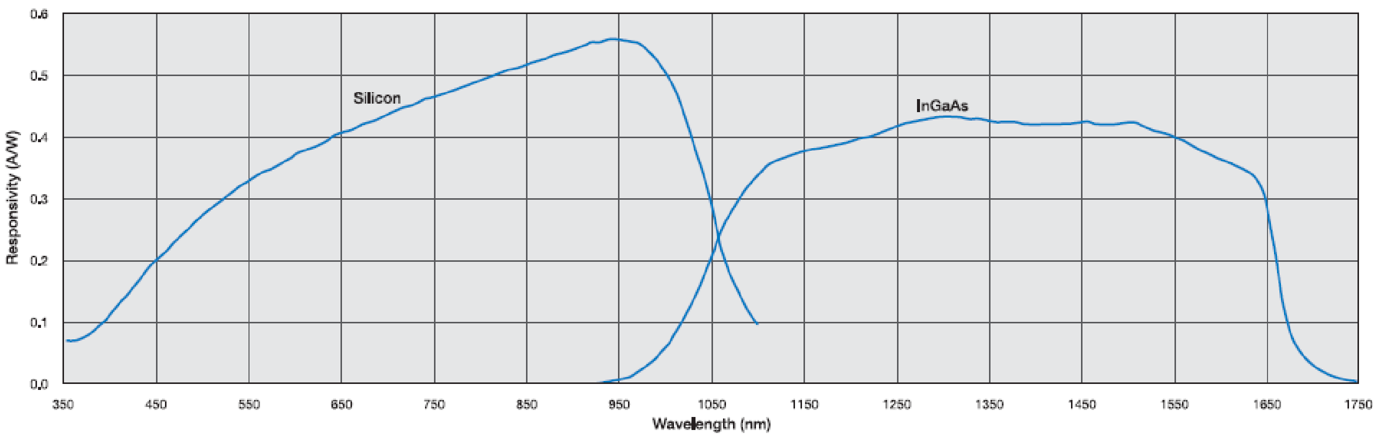

K cutout in paper

front-lit
visible

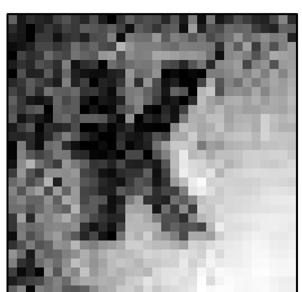

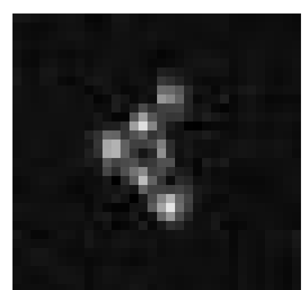

back-lit
IR LEDs

CS Low-Light Imaging with PMT

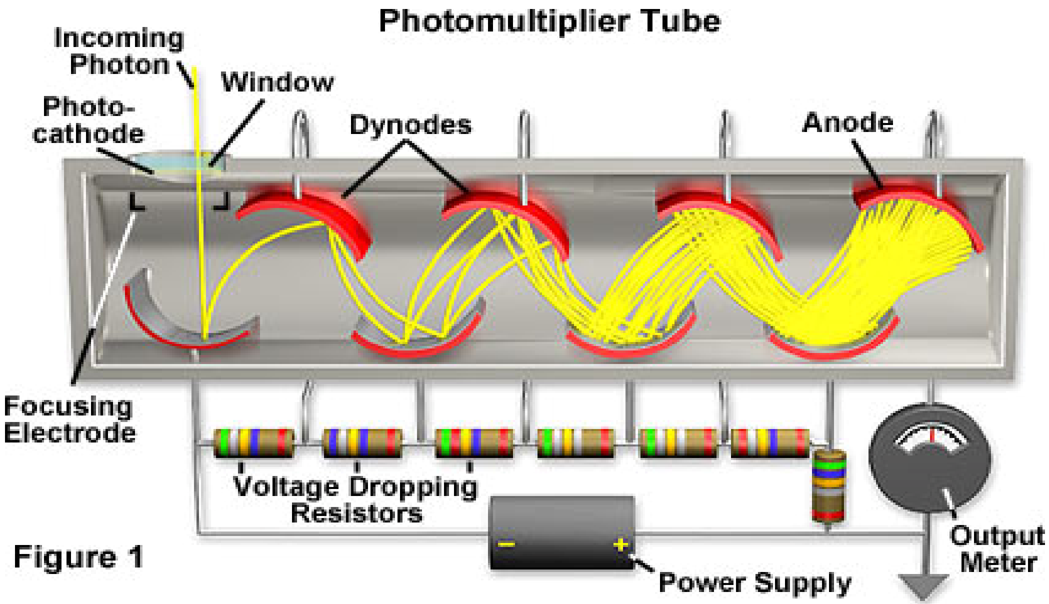

true color low-light imaging

256 x 256 image with 10:1 compression

[Nature Photonics, April 2007]

Video Acquisition

- Measure via stream of random projections
  - shutterless camera
- Reconstruct using sparse model for video structure
  - 3-D wavelets (localized in space-time)

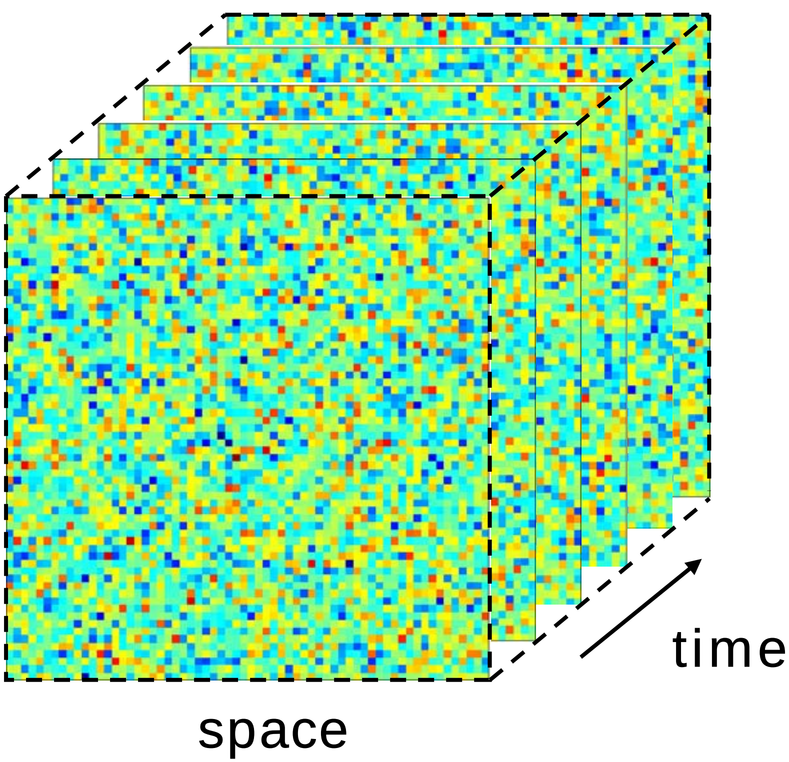

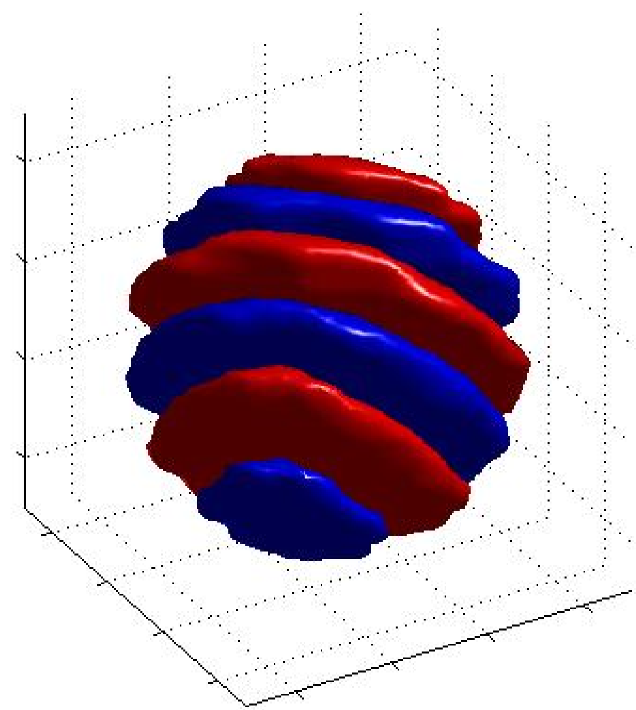

3D complex wavelet

original 64x64x64

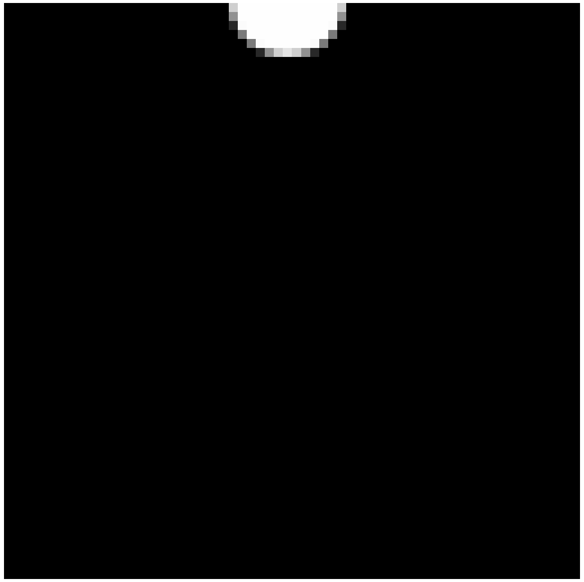

3-D wavelet thresholding

2000 coeffs, MSE = 3.5

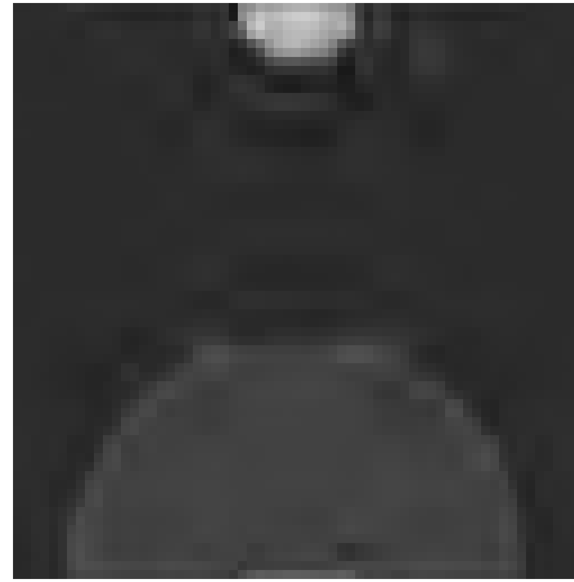

frame-by-frame 2-D CS recon

20000 coeffs, MSE = 18.4

joint 3-D CS recon

20000 coeffs, MSE = 3.2

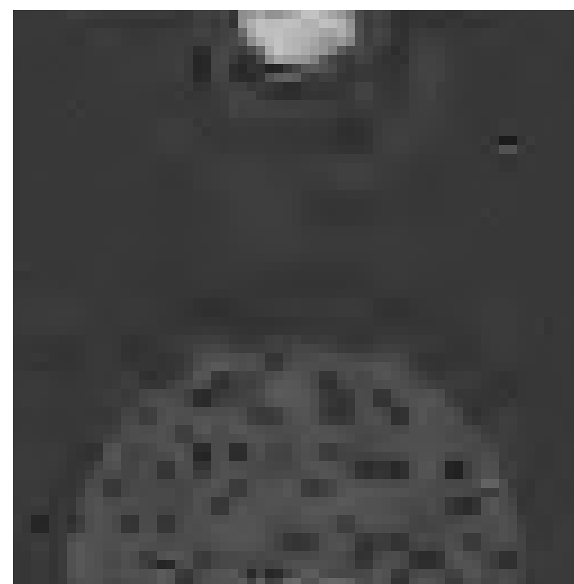

Miniature DMD-based Cameras

- TI DLP "picoprojector" destined for cell phones

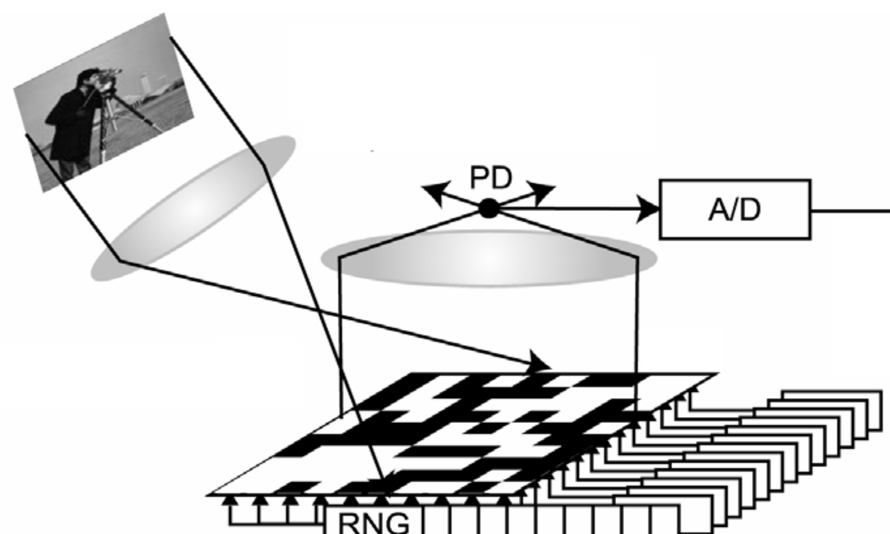

Slashdot

News for Nerds. Stuff that matters.

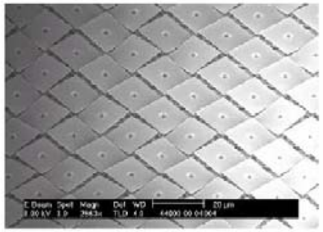

Oops, crash, seven million years of bad luck!?!

This is me skydiving

.

This is me swimming with dolphins

.

This is me at the Grand Canyon

.

Contorno

Information Scalability

Information Scalability

- **Random projections ~ sufficient statistics**

- Same random projections / hardware can be used for a range of different signal processing tasks
  reconstruction > estimation > recognition > detection

- Many fewer measurements may be required to detect/classify/recognize than to reconstruct

- Example applications:
  - adaptive cameras
  - smashed filter: compressive matched filter
  - non-imaging cameras
  - meta-analysis

Attentive CS Video Camera

[Ilan Goodman, Don Johnson]

- Detect activity from random measurements
- Detection requires far fewer measurements than reconstruction
  - 320x240 pixels x 24 bits/pixel x 20 frames per second
    = **36,864,000 bits per second**
  - *detect activity* from statistics of
    6 CS measurements/second x 4 bits/measurement
    = **24 bits/second**

red = rate throttled back

Hyperspectral Image Classification

- 3D random projections of hyperspectral data cube
- Classify/segment rather than reconstruct

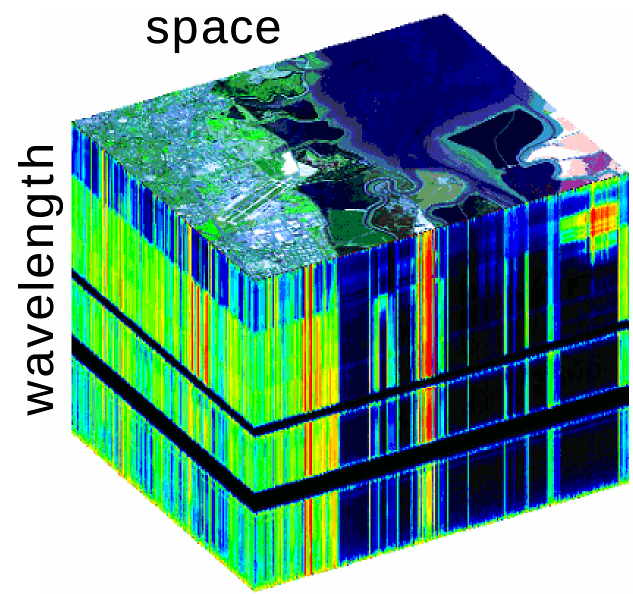

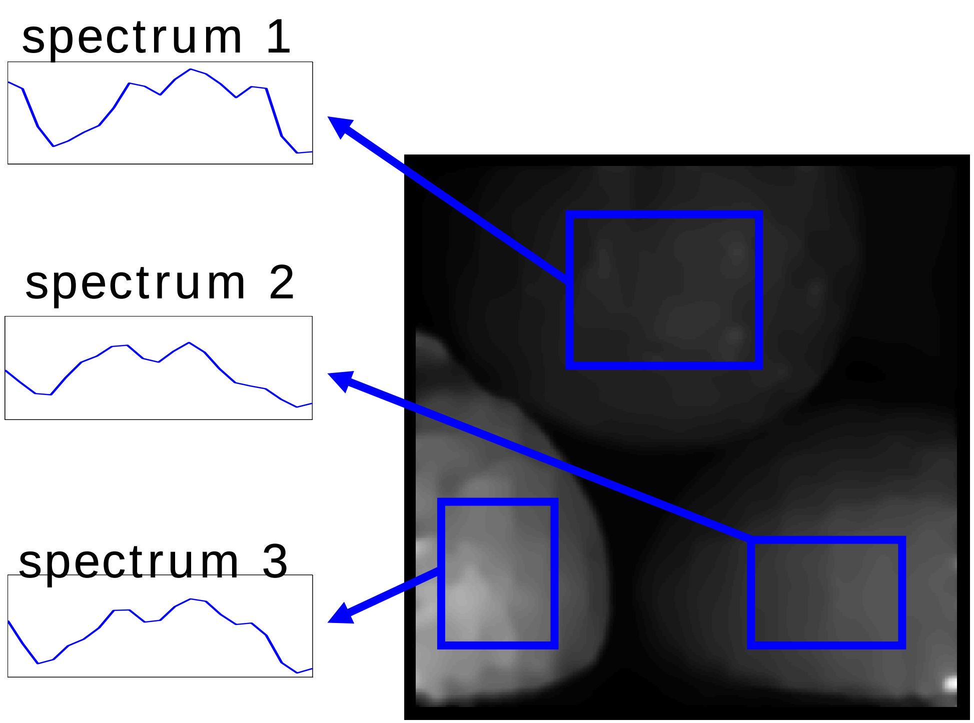

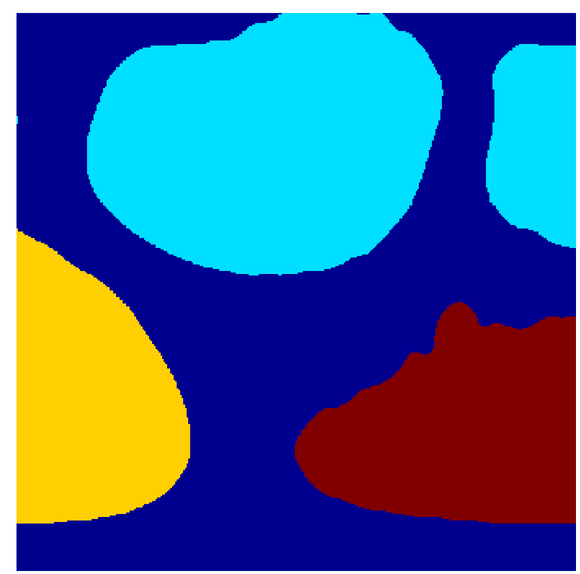

Matched Filter

- For signal classification when templates are parametrically **transformed**
  - ex: shift/rotate/scale
  - formulated via GLRT
- Underlying geometry: low-dimensional **manifold**
- Classification: nearest manifold search

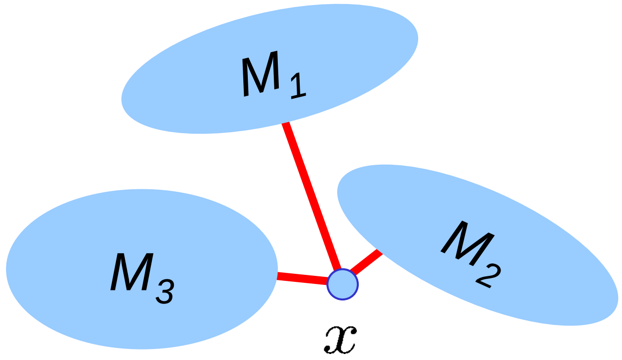

Smashed Filter

- Dimension-reduced GLRT for parametrically transformed signals
- Key theoretical ingredient: *manifold structure preserved by random projections*
- Classification: nearest manifold search

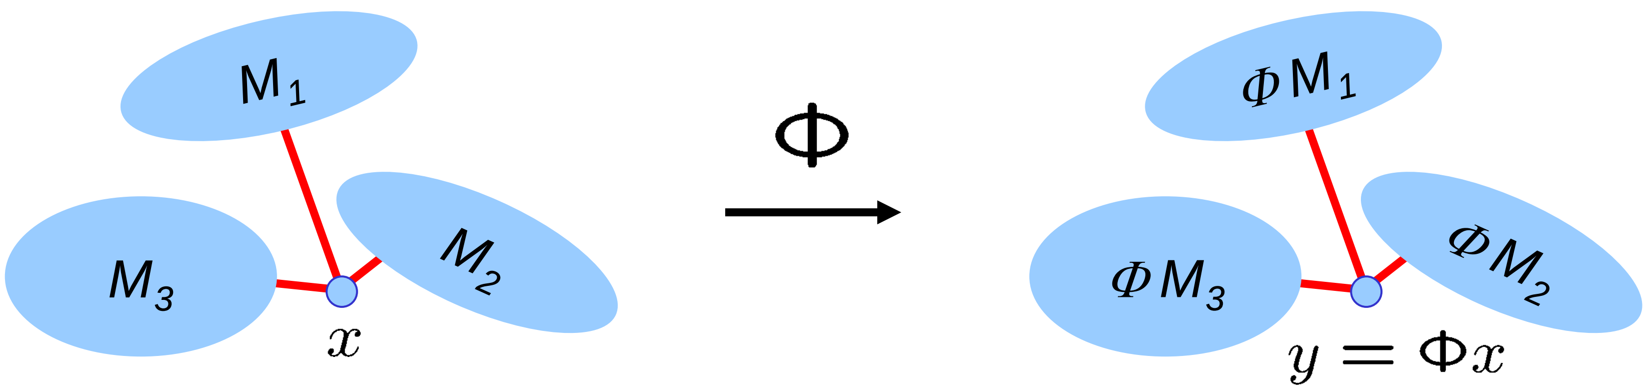

Smashed Filter – Experiments

- 3 image classes: tank, school bus, SUV
- N = 65536 pixels
- Imaged using single-pixel CS camera with
  - unknown shift
  - unknown rotation

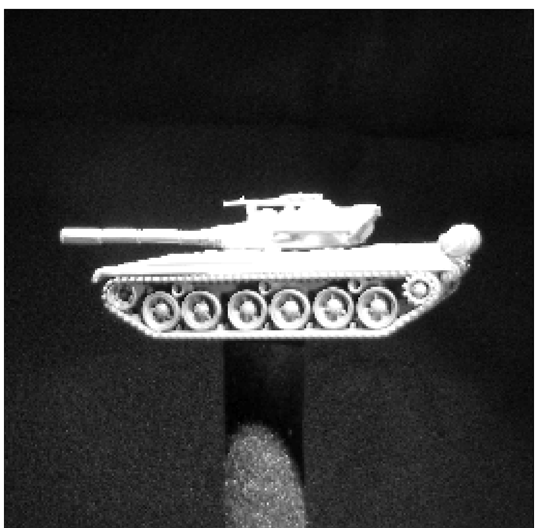

Smashed Filter – Unknown Position

- Object shifted at random (K=2 manifold)
- Noise added to measurements
- Goal: identify most likely position for each image class
  identify most likely class using nearest-neighbor test

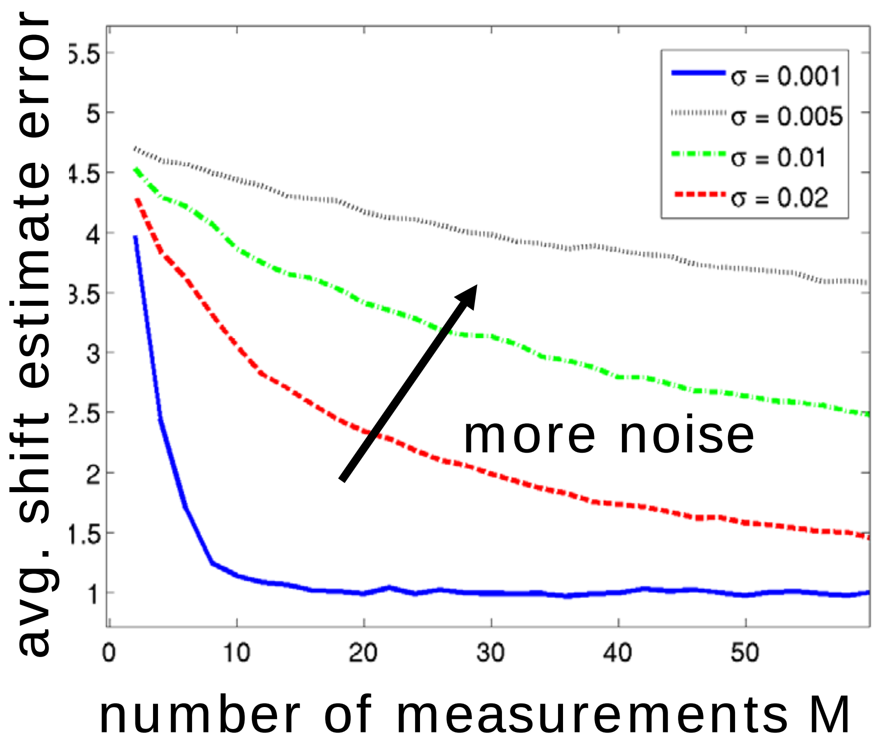

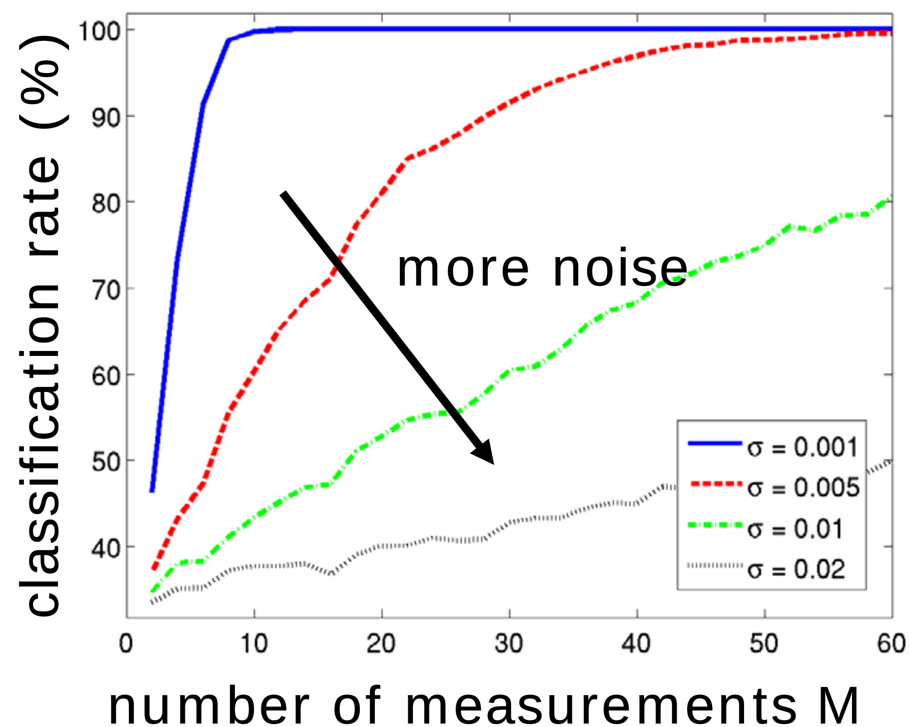

Smashed Filter – Unknown Rotation

- Object rotated each 2 degrees
- Goals: identify most likely rotation for each image class
  identify most likely class using nearest-neighbor test
- Perfect classification with as few as 6 measurements
- Good estimates of rotation with under 10 measurements

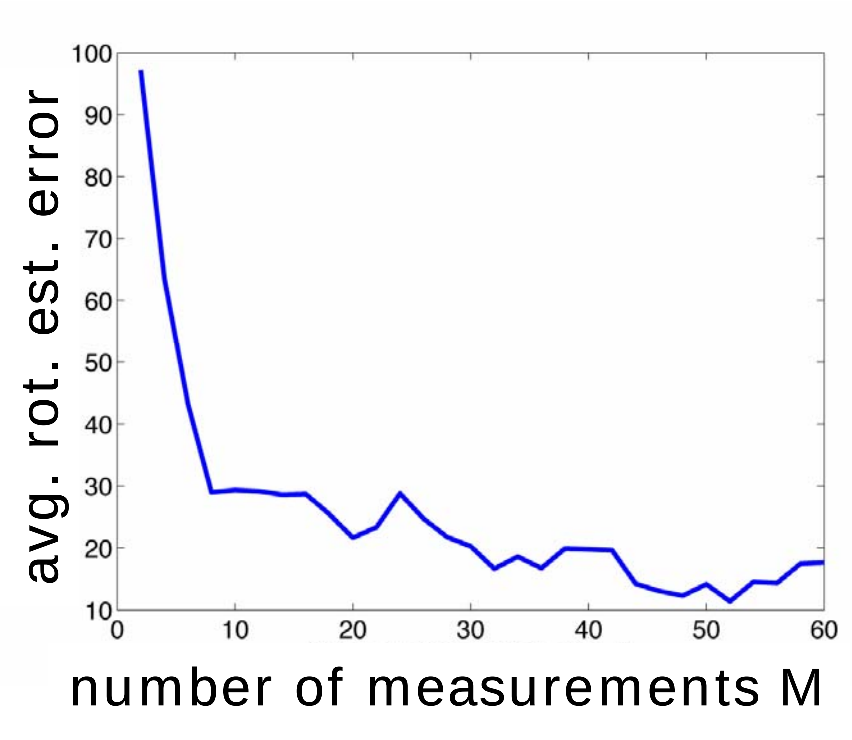

Dolce

Other Compressive
Camera Architectures

Random Lens Camera

- Computes random sums using random mirror
- Use regular CCD array to acquire many random sums at once
- CS reconstruction yields super-resolved image

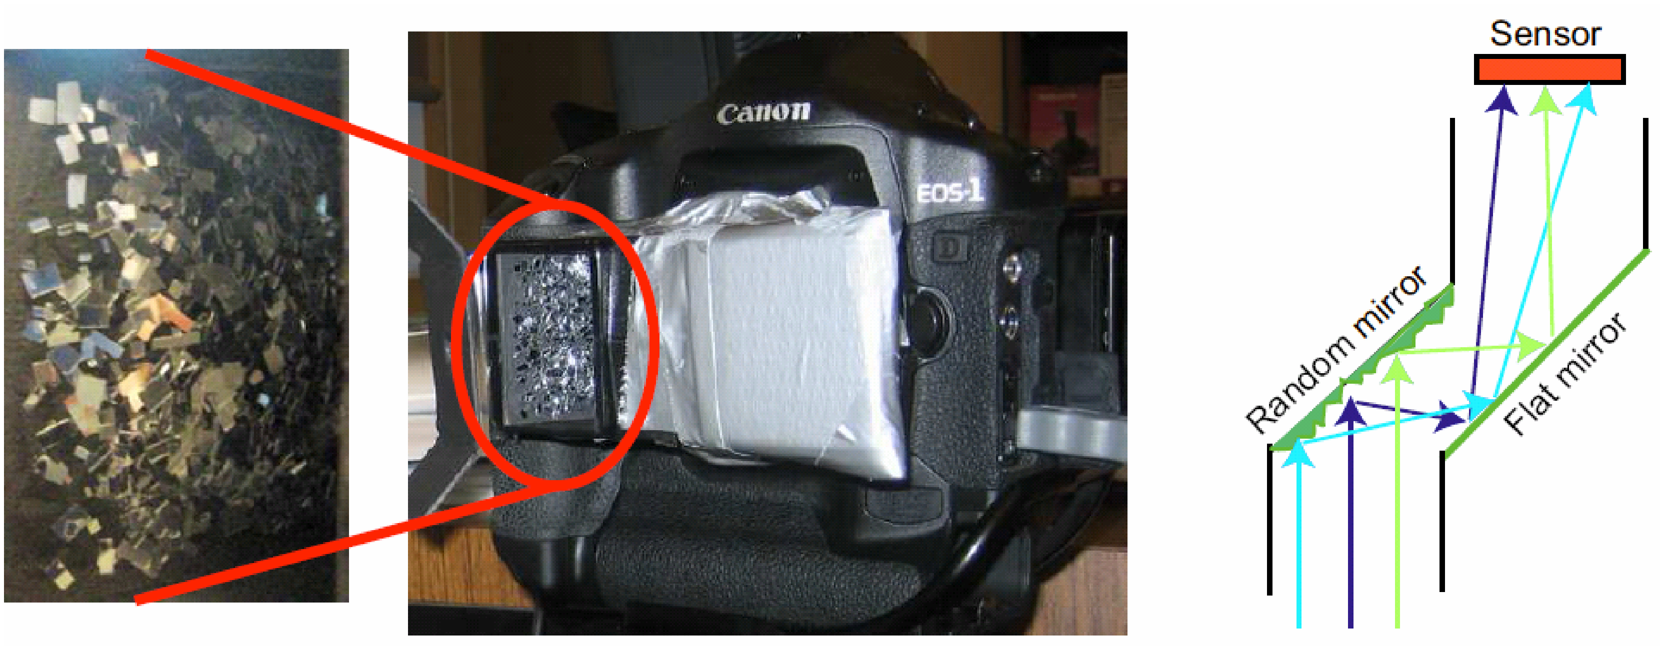

[Rob Fergus, Antonio Torralba, Bill Freeman @ MIT]

Thin Cameras

- Dave Brady @ Duke
- Thin cameras < 2.5mm
- Based on coded aperture, lenslets

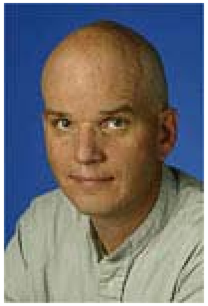

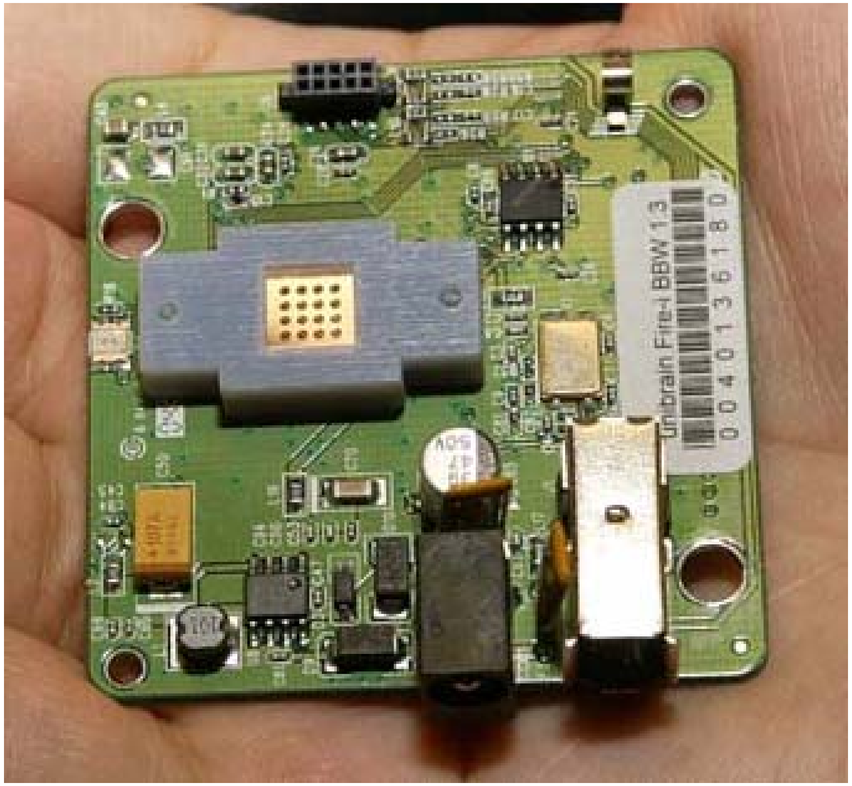

Café

Conclusions

What's In it for You?

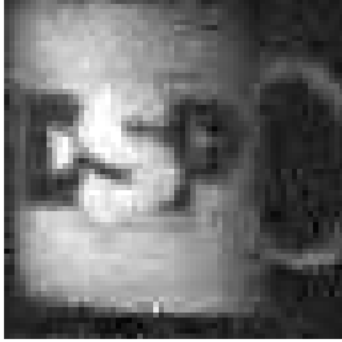

- **Compressive sensing**
  - exploits signal sparsity/compressibility information
  - based on new uncertainty principles
  - Sudoku-like reconstruction from random measurements
  - integrates sensing, compression, processing
  - enables new sensing architectures and modalities
  - most useful when measurements are expensive

- CS measurements are **information scalable**

  reconstruction > estimation > classification > detection

- Selected mid/long-term **applications**
  - cameras and imagers where CCDs and CMOS imagers are blind (science, military)
  - security applications (potential for low cost / low power)
  - large camera arrays (compressibility gain with multiple cameras)
  - advanced algorithms for today's cameras (eg: deblurring)

dsp.rice.edu/cs

Rice University DSP

Contact

Richard Baraniuk and ***Kevin Kelly***

ECE Department

Rice University

richb@rice.edu, kkelly@rice.edu

dsp.rice.edu/cs

dsp.rice.edu/cscamera