Computational, Compressive, and Task-Specific Imaging

Mark A. Neifeld
University of Arizona
College of Optics and ECE Department

OUTLINE

1. History and Motivation
2. Reconstructive Imaging
3. Task-Specific Imaging
4. X-ray Imaging
5. RF Imaging

PART 1
History and Motivation

Camera Evolution

Camera History Through 1990

- 400 BC — camera obscura (pinhole)
- 1600 AD — camera obscura (lens)
- 1830 AD — plate recording (Daguerre)
- 1890 AD — film recording (Eastman)
- 1961 AD — digital concept (Lally)
- 1975 AD — digital prototype (Kodak)
- 1988 AD — JPEG standard
- 1990 AD — digital camera with PC connection

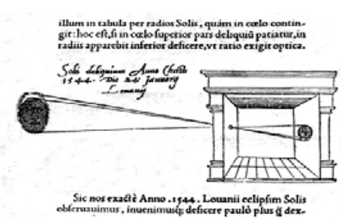
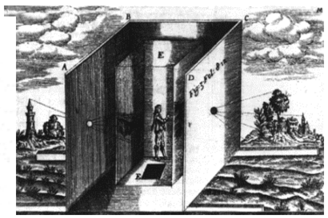

Human measurement


Analog measurement

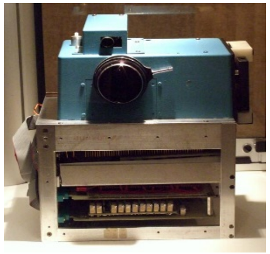


Digital measurement

Camera Status in 1990

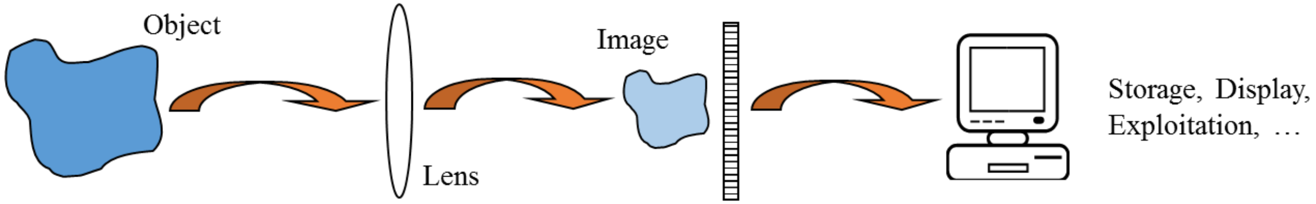

For almost 2400 years cameras have measured pretty pictures.
This approach is no longer necessary or desirable.

Camera Revolution

Camera History 1995-2015

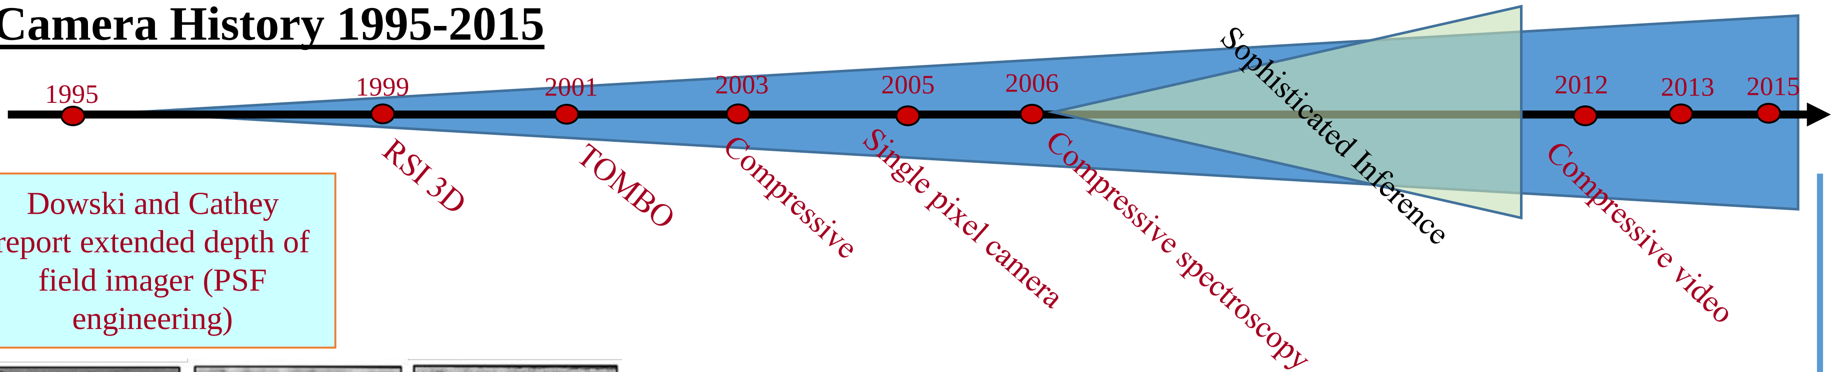

Dowski and Cathey report extended depth of field imager (PSF engineering)

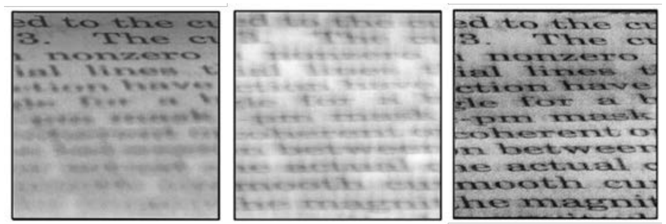

conventional | cubic raw | reconstruction

Computational Cameras in 2015

PSF Engineering (EDOF, super-resolution, aberration control, …)
Compressive 3D (e.g., light-field, x-ray CT and/or ladar)
Compressive RF (e.g., radar imaging, tracking, AIT, …)
Multi-spectral imaging/video
Multi-domain optimization and task-specific variations

Camera Status in 2015

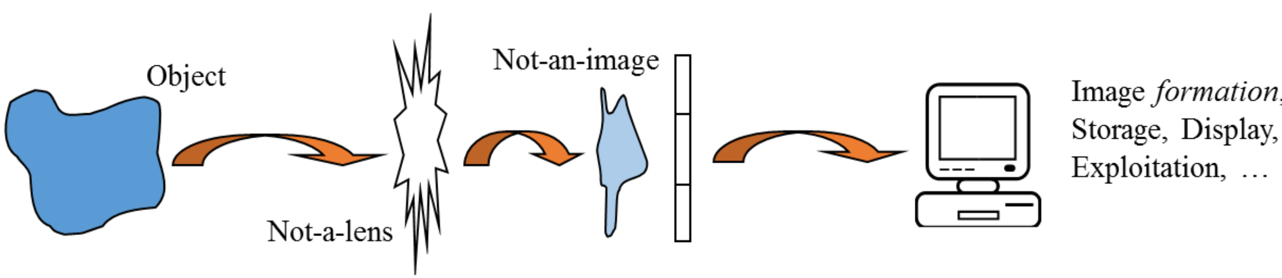

The revolution: computation can/should be central to image formation

Philosophical Underpinnings of Compressive Imaging

Images are Redundant

- This is true for all modalities and applications.
- It is true because the world is highly correlated in space, time, and wavelength.
- The world is made of objects … not pixels.


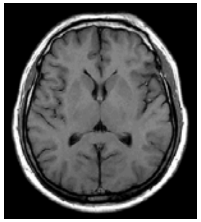
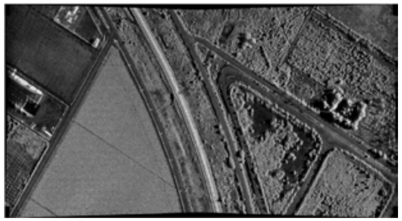

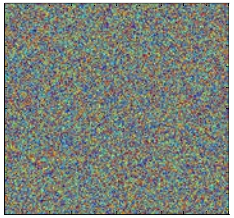

This Should Make Us Uncomfortable

- Practically – we are spending resources on measuring redundant data.
- Theoretically – we know more than band-limited so we need to “fix” sampling theorem.
- Cybernetically – sometimes an image is never intended for human consumption.

A Practical Perspective

Measurement is the application of resources toward the extraction of important information from a physical system

Information Bottleneck

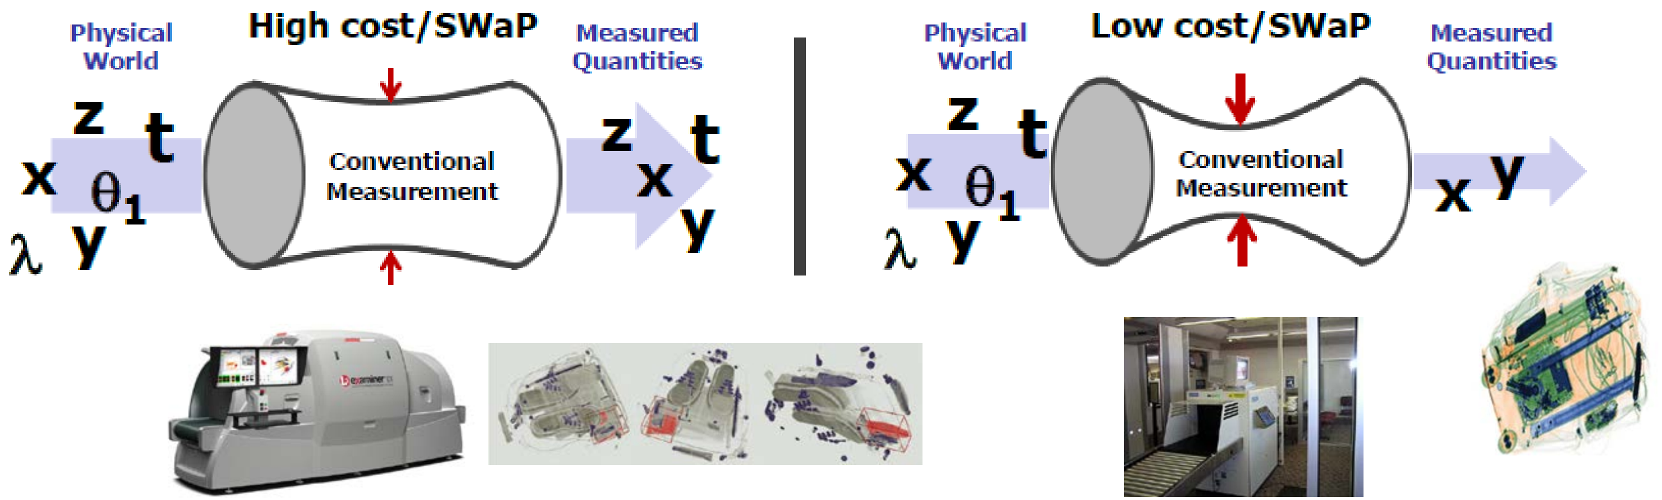

Motivational Observations:

- Measurement is the only means to extract information from the world.
- Measurements have cost (size, weight, power, latency, …).
- The physical world offers more variables than we can afford to measure.
- Bottleneck demands careful selection of measurements that convey most useful information.

Important to think of measurement in terms of resource costs

A Theoretical Perspective

- Images are redundant – this is why image compression is so effective
- Redundancy can be viewed as sparsity in some transform domain (L of N nonzero coefficients)

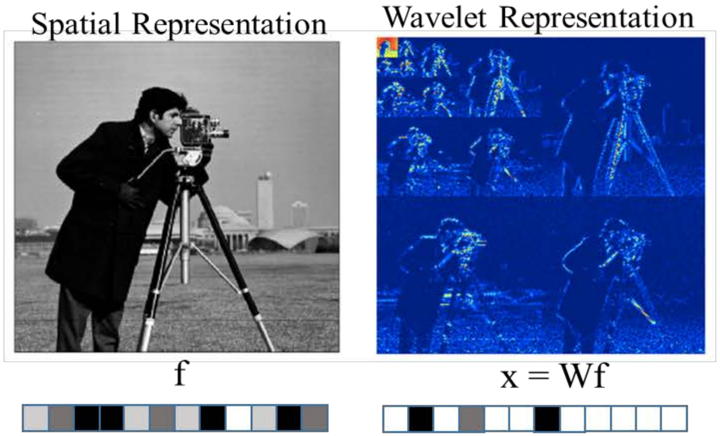

- How can measurements on f leverage this knowledge?
- Traditional measurement requires O(N) samples.
- Donoho proved that compressive measurements using random projections will require only O(L)

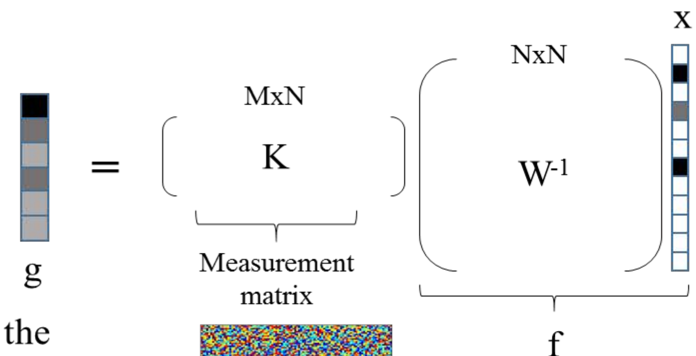

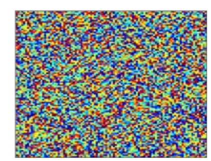

- Reconstruct by solving an inference problem – what is the best estimate of x given (a) what I've measured, (b) the requirement for sparsity, and (c) anything else I know about x (e.g., positivity, prior distribution, …) ?

$$x = \text{argmin} \{ |x|_1 + \lambda |g-KW^{-1}x|^2 + \gamma P(x) \}$$

Sparsity   Data Agreement   Regularizer

Nyquist was overly pessimistic. Stronger prior knowledge than band-limited is generally available

Measurement cost now scales with information dimension not signal dimension

The Weighing Problem

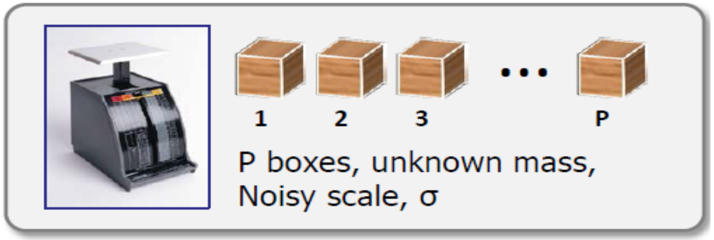

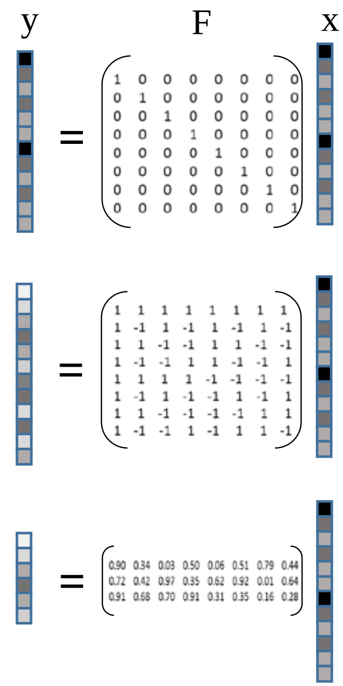

Sequential procedure

- P measurements needed
- SNR = SNR_0
- No inversion required

Multiplexed procedure

- P measurements needed
- SNR = (P/2) SNR_0
- Linear inversion required

Compressive procedure

- P boxes, S have non-zero weights
- O(S) measurements needed
- SNR = (S/2) SNR_0
- Nonlinear inversion required

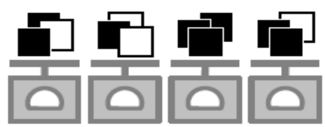

Question: What are the optimal combinations?
Answer: It depends upon priors, task, and measurement physics/noise

DARPA KECOM

Knowledge Enhanced Compressive Measurement (KECoM) relaxes the information bottleneck via measurement optimization

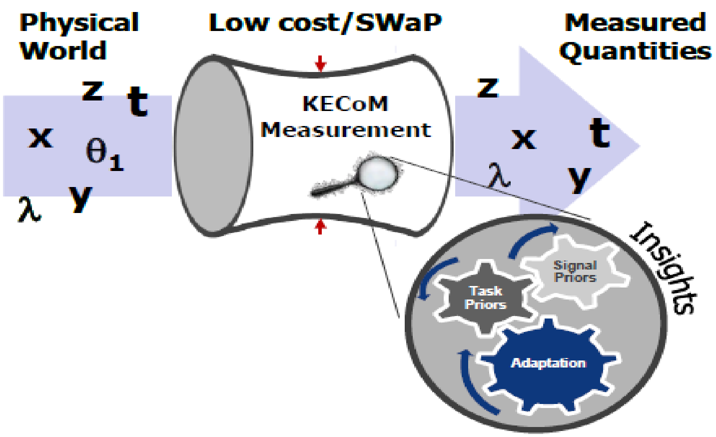

KECoM Insights:

1. Don't measure what you already know (i.e., via priors and/or previous measurements)
2. Don't measure what you don't need to know (i.e., measurements should be task specific)

Approach and Benefits

- Minimize measurement resources via ***compressive sensing***
- Optimize measurement choices using ***task-specific information***
- Example EO/IR result shows 30x SNR advantage with 10x fewer measurements compared with conventional imaging:
  - Increase Pd and/or decrease Pfa
  - Reduce measurement time/energy
  - Reduce deployment and/or operation costs
  - Explore complementary measurement modalities
  - Improve material specificity and reduce divestiture

Applications of Interest:

1. **VIS/IR Imaging** (e.g., UAV, soldier borne, and other mobile applications)
2. **RF Imaging** (e.g., sparse arrays, waveform design, under-sampled ADC.
3. **Spectroscopy** (e.g., standoff chemical sensing)
4. **SIGINT** (e.g., cooperative and uncooperative communications)
5. **LADAR** (e.g., remote 3D imaging)
6. **X-Ray Imaging/Threat Detection** (e.g., airport security)

The CS/TSI advantage can be "sliced" in a variety of ways depending upon application needs

PART 2
Reconstructive Imaging (VIS/IR)

Motivational Observations

1. Why are images typically measured as collections of pixels?
   - History = Original consumers were humans who wanted pretty pictures.
   - Physics = Image-formation using glass (or a pinhole) is "straightforward."
   - Technology = Previous lack of electronic detection/post-processing.
   - Mind Set = If I can't "see it" then it isn't there.

2. What is the first thing we do after we measure 100 Mpixel image?
   - Use compression to throw away the redundant parts.
   - Maybe we can push some compression into the measurement domain.

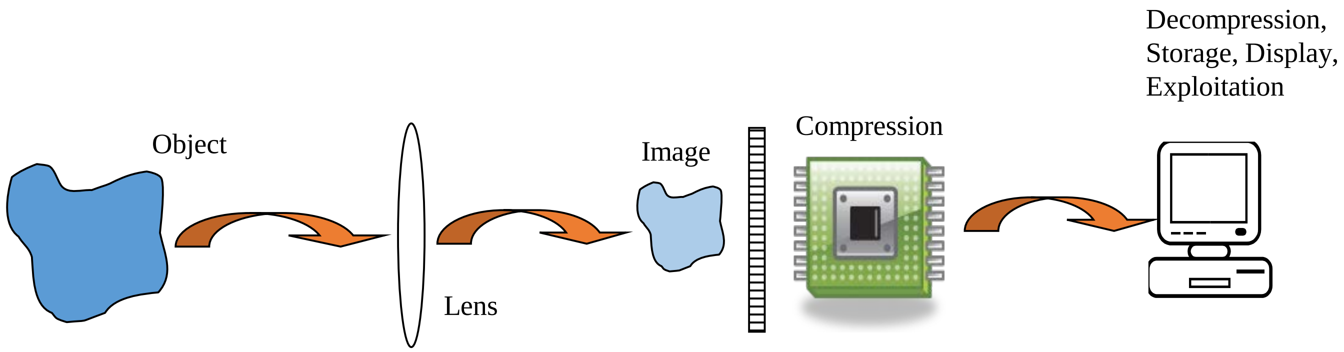

Motivational Observations

1. Why are images typically measured as collections of pixels?
   - History = Original consumers were humans who wanted pretty pictures.
   - Physics = Image-formation using glass (or a pinhole) is “straightforward.”
   - Technology = Previous lack of electronic detection/post-processing.
   - Mind Set = If I can’t “see it” then it isn’t there.

2. What is the first thing we do after we measure 100 Mpixel image?
   - Use compression to throw away the redundant parts.
   - Maybe we can push some compression into the measurement domain.

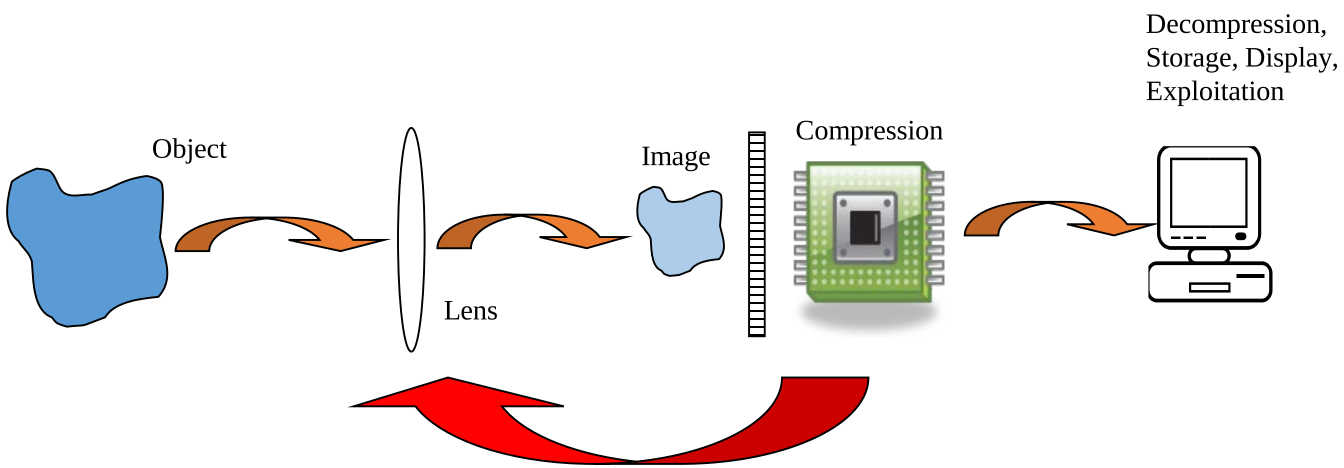

Can compression be combined with collection?

Feature-Specific Compressive Imaging

- Applications in the MWIR and LWIR can be dominated by the cost of focal planes.
- Compressive imaging can substantially reduce FPA costs with little reduction in image quality.

Conventional

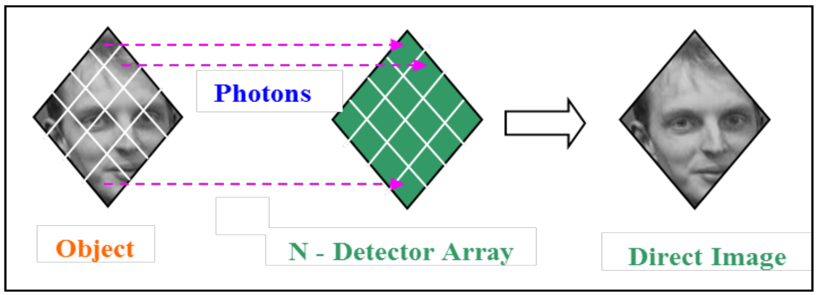

Feature-Specific

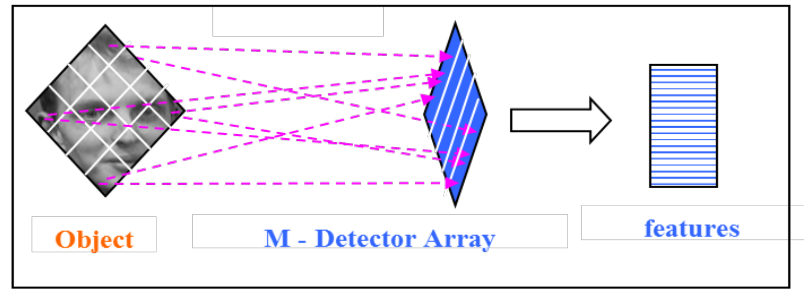

- Conventional imagers measure a large number (N) of pixels
- Compressive imagers measure a small number (M<<N) of features
- Features are simply projections $y_i = (\mathbf{x} \cdot \mathbf{f}_i)$ for $i = 1, \ldots, M$
- Benefits of projective/compressive measurements include
  - increased measurement SNR → improved image fidelity
  - more informative measurements → reduced sensor power and bandwidth
  - enable task-specific imager deployment → information optimal
- Previous feature-specific imaging: Neifeld (2003), Brady (2005), Donoho (2004), Baraniuk/Kelly (2005)

PCA, ICA, Wavelets, Fisher, multi-spectral …

DCT, Hadamard, …

random projections

Example Architectures

Sequential Compressive Imaging

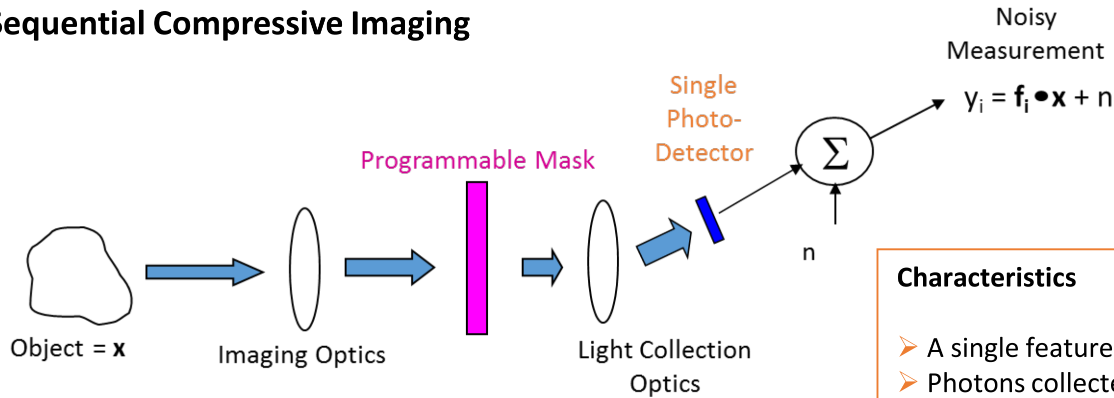

Characteristics

- A single feature is measured in each time step (noise BW ~ M/T)
- Photons collected on a single detector (measured signal ~ 1/M)
- Unnecessary photons discarded in each time step (1/2)
- Reconstruction computed via post-processing

Parallel Compressive Imaging

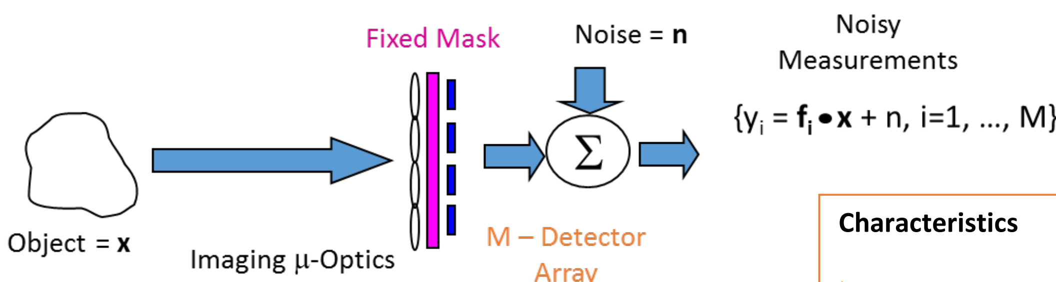

Characteristics

- All M features are measured in a single time step (noise BW ~ 1/T)
- Photons collected on M << N detectors (measured signal ~ 1/M)
- Unnecessary photons discarded in each channel (1/2)
- Reconstruction computed via post-processing

Kernels are completely arbitrary beyond physical requirement for positivity and energy conservation

Common Compressive Imaging Tradeoff

- SNR determined by number of object photons
- Fair comparison requires equal numbers of photons
- All feature measurements must share photon budget ➔ $|P_M| = 1$
- Results based on parallel architecture

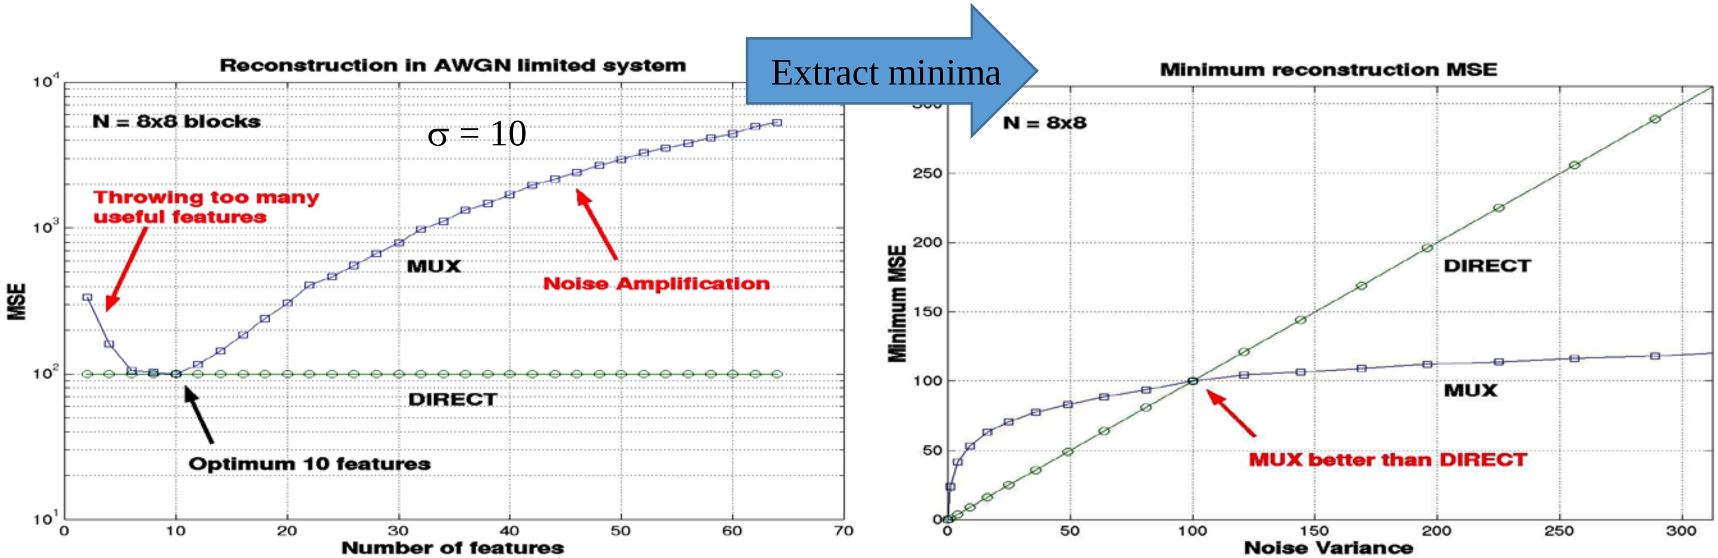

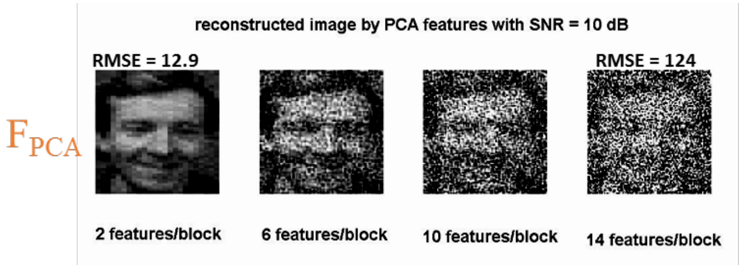

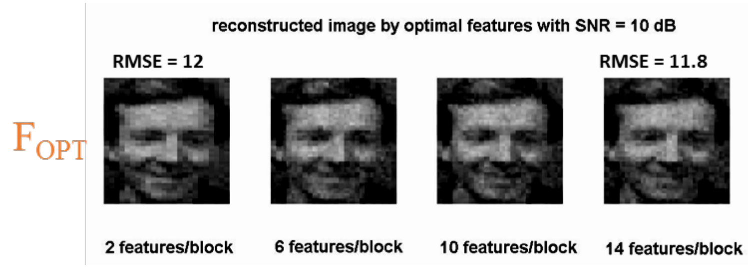

Compressive Imaging via PSF Engineering

PSF Engineering for sub-pixel resolution:

Existing solutions trade field of view and resolution!

Q: How can PSF engineering improve UAV cameras?
A: Provide optics with additional degrees of freedom

- Consider use of extended point spread function
- Design issue #1: retain full optical bandwidth
- Design issue #2: tradeoff SNR for condition number
- Pseudo-Random Phase masks for extended PSF

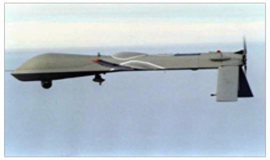

Soda Straw

Poor Resolution

Problem = pixel-limited resolution

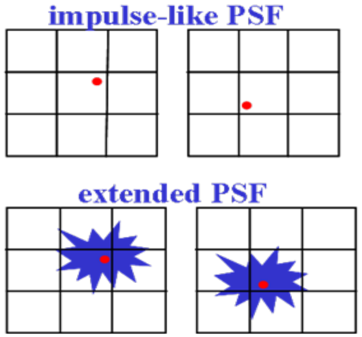

Pseudo-Random Phase mask Enhanced Lenslet

PSF Engineering

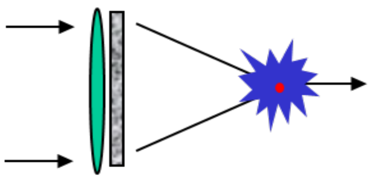

Modulation Transfer Function

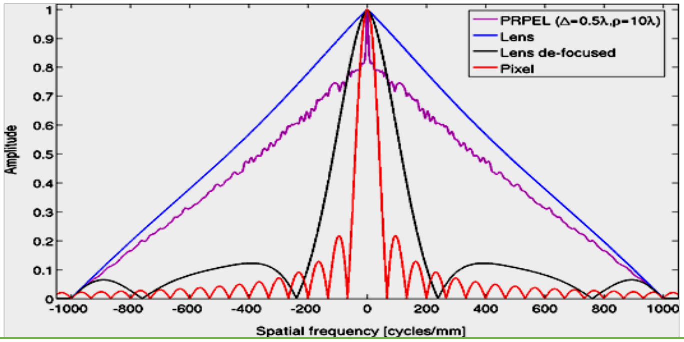

PSF Engineered Super-Resolution

Sine Phase mask Enhanced Lenslet $\phi(x)=\sum_{i=1}^{N}\alpha_i \sin(\omega_i x+\theta_i)$ α

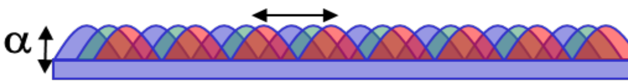

- Pick N=3: yields 12 free parameters for optimization.
- Optimization Criterion = RMSE
- RMSE computer over objects class using LMMSE operator

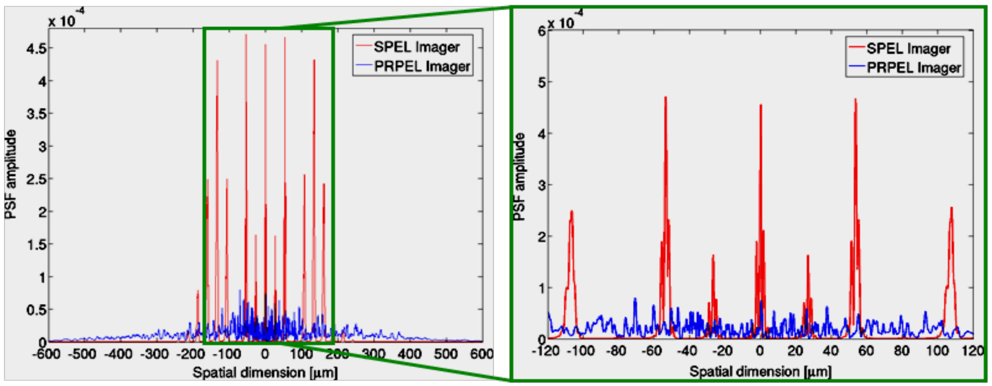

Observations

- Note smaller support of SPEL PSF compared to PRPEL PSF.
- SPEL PSF has more efficient photon-distribution.
- SPEL PSF also contains sub-pixel structure.

Conventional

PRPEL-LIN

PRPEL-RL

Capability enabled NOT by simple image processing … but by joint design (MDO).

Joint design of optics and post-processing overcomes pixel-limited resolution

PART 3

Task-Specific Imaging

Image Information Content

512 × 512 × 3 × 8 = **6.2 Mb**

64 × 64 × 1 × 8 = **32Kb**

Compression

2.1 Mb

24 Kb

- Information content source requires probability density ρ(**r**)

$$\text{Shannon Entropy}: \; J = -\int \rho(\mathrm{r}) \log \rho(\mathrm{r}) d^n r$$

PROBLEM: In general ρ(**r**) is very complex/unknown and high-dimensional

Task Specific Information Concept

- Information content is *task specific.*

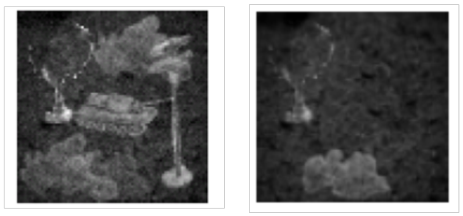

Detection task:

Probability of presence/absence = ½
Information content < **1** bit

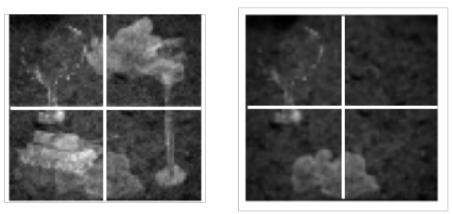

Detection and Localization task:

Probability of tank being absent = ½
Probability of occurrence in each region = ⅛
Information content < **2** bits

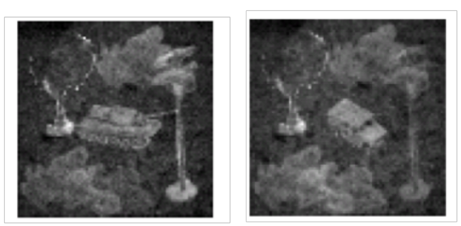

Classification task:
Probability of each target the = ½
Information content < **1** bit

- How do we ***quantify*** the ***task specific information*** (TSI)?

Task Specific Source Encoding

Virtual source → X → Encoding → Y = *C*(X) →

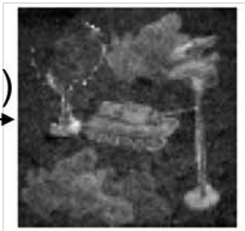

SCENE

C(X) stochastically encodes X and produces scene Y

- Detection task: *presence/absence of target is of interest*
  - Virtual source variable X must be binary.
  - X = 1 or 0 implies tank present or absent.

X = 1 (Tank present)

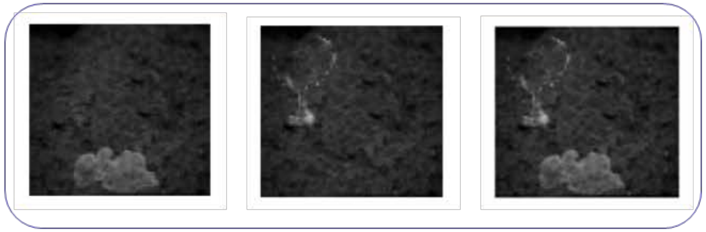

X = 0 (Tank absent)

An image may be viewed as an encoding of task-specific information

Task Specific Information Definition

Imaging chain block diagram

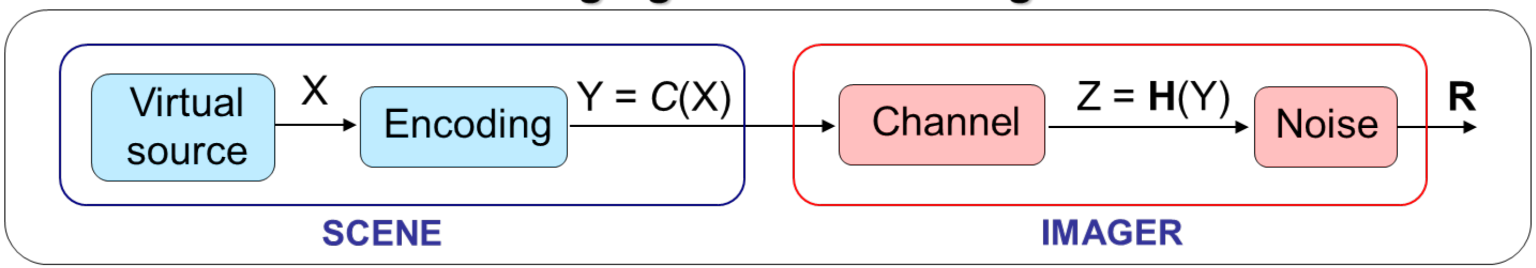

- Imager is characterized by channel **H** and noise *n*
- Definition for Task Specific Information:

$$\boxed{TSI \equiv I(X;R) \leq J(X)}$$

Mutual information between X and *R*

Always bounded by the entropy of X

Entropy *J*(X) → maximum task specific information content

$$\text{TSI} = H(R) - H(R|X) = H(X) - H(X|R)$$

TSI may be discrete (i.e., previous examples) or continuous

2-Class Detection Problem

- Detection task: *presence/absence of target is of interest*
  - Virtual source variable X must be binary.
  - X = 1 or 0 implies tank present or absent.

X = 1 (Tank present)

X = 0 (Tank absent)

Sequential Architecture

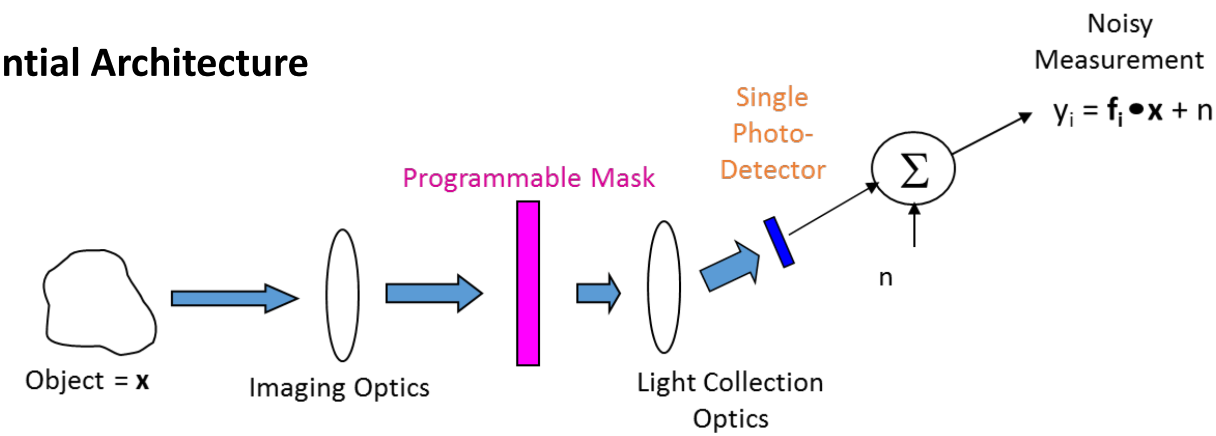

- Dual-rail is employed to realize negative quantities
- Energy conservation is enforced via the photon count constraint
- Only AWGN is considered in the results to follow

TSI Comparison Among Compressive Imagers

- Binary detection problem (tank present/absent) → 1 bit of TSI.
- Include positivity and energy conservation constraints.
- Compare various known measurement matrices with conventional imager.
- TSI optimization only over photon allocation (e.g., dwell time).

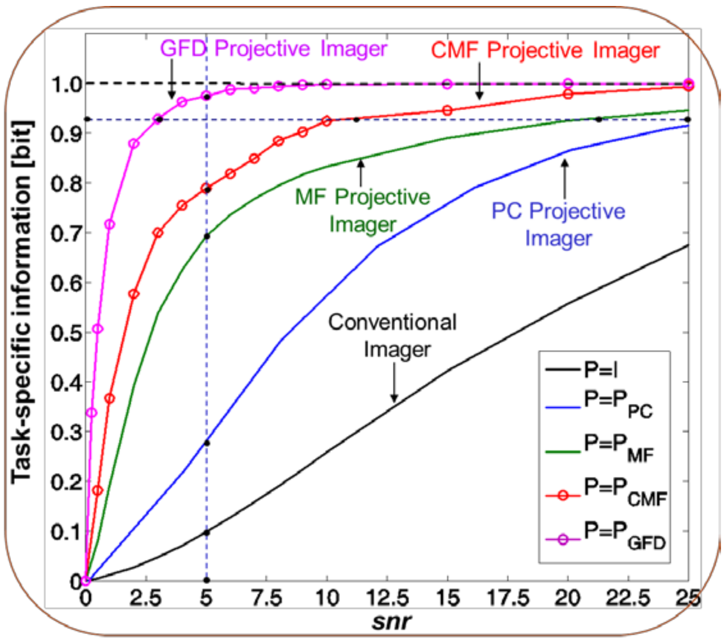

Fano Inequality

$$TSI(X;R) - J(X) = J(X|R) \leq P_e \log(\aleph-1) + J(P_e)$$

$P_e=10^{-2}$ ⟺ TSI=0.92 bit

P = I
snr = 5 → TSI = **0.10** bit
P = P_{PC}
snr = 5 → TSI = **0.28** bit
P = P_{MF}
snr = 5 → TSI = **0.69** bit
P = P_{CMF}
snr = 5 → TSI = **0.79** bit
P = P_{GFD}
snr = 5 → TSI = **0.96** bit

P = I
TSI = 0.92 bit @ *snr* = **44**
P = P_{PC}
TSI = 0.92 bit @ *snr* = **28**
P = P_{MF}
TSI = 0.92 bit @ *snr* = **21**
P = P_{CMF}
TSI = 0.92 bit @ *snr* = **11**
P = P_{GFD}
TSI = 0.92 bit @ *snr* = **3**

Even sub-optimal compressive measurement can provide 14.7x SNR improvement relative to conventional imaging

TSI Optimal Extension

- Binary detection problem (tank present/absent) → 1 bit of TSI.
- Include positivity and energy conservation constraints.
- Compare various known measurement matrices with conventional imager.
- TSI optimization yields fully optimal measurement vectors.

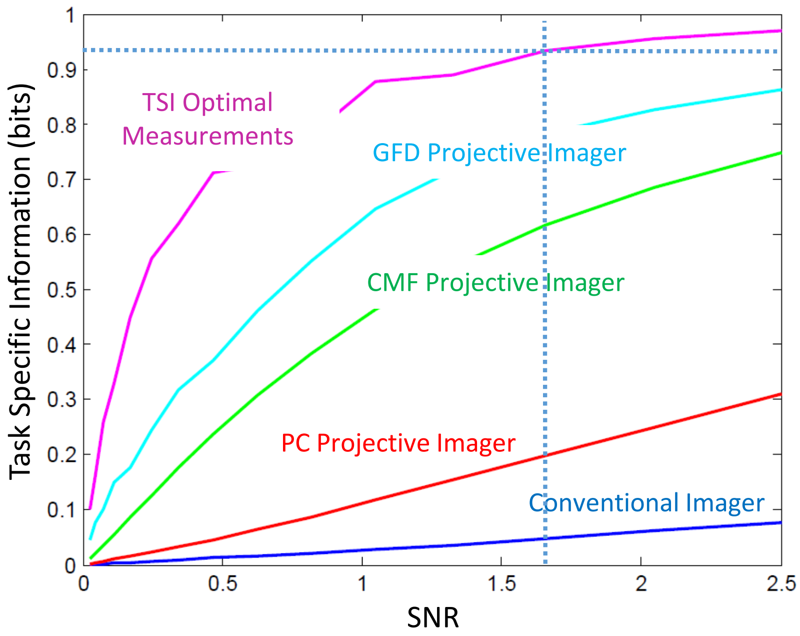

Information optimal compressive measurement can provide 27.5x SNR improvement (i.e., to achieve Pe=0.01) relative to conventional imaging

RMSE Benefit of TSI-Optimal Reconstruction

- Results for blockwise compressive imaging using 4x4 blocks and measurement SNR = 30dB.
- Conservation of energy via photon count constraint.
- Correlated designs provide 37% RMSE improvement over uninformed designs at 4x compression
- Correlated designs provide >2x compression advantage at RMSE = 4%

Image Quality versus Number of Measurements

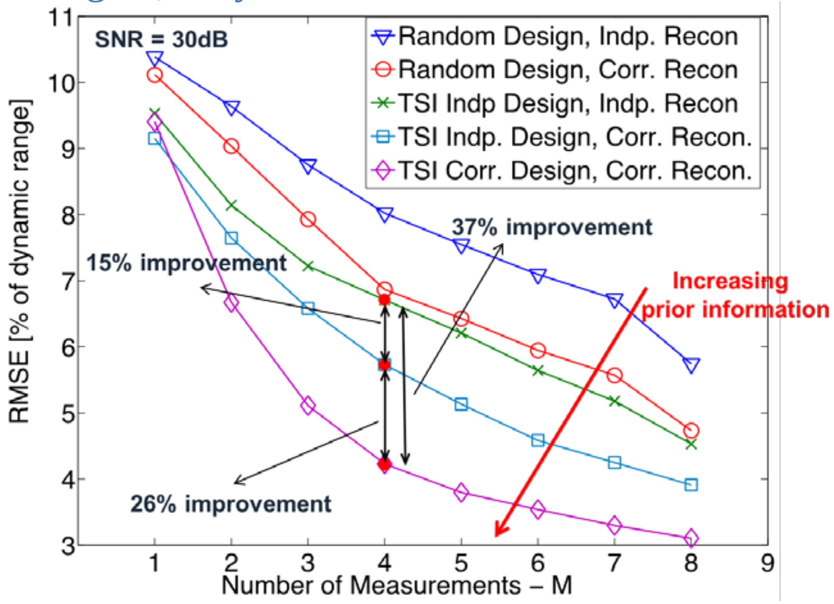

4x compression

Random Design
Correlated Reconstruction

Correlated Design
Correlated Reconstruction

PART 4
X-Ray Imaging

Fixed-Gantry X-Ray Tomography

- Airport (and other infrastructure) security relies on x-ray imaging for threat detection
- Current systems combine automated cueing with operator-in-the-loop detection
- Rotating gantry systems are common
- Fixed-gantry systems are emerging due to their improved speed and flexibility.

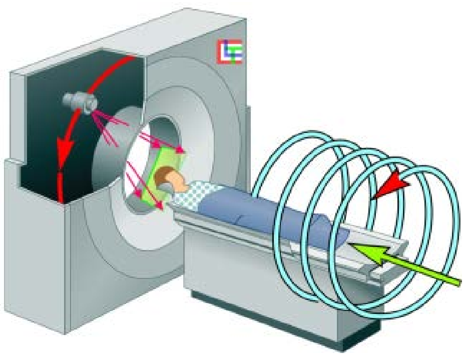

Rotating Gantry

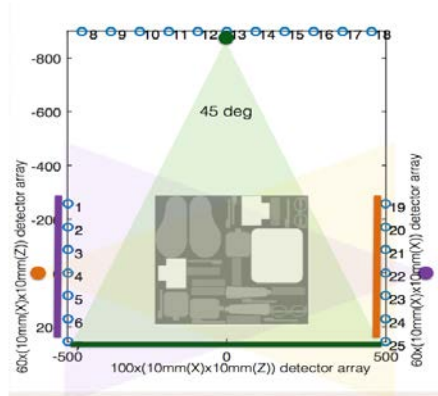

Fixed Gantry

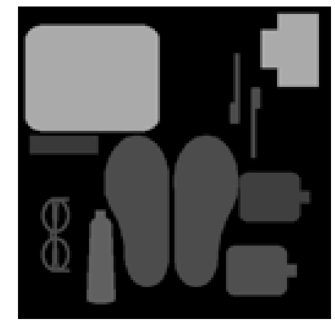

Full Rank Reconstruction

- Architecture of interest has ~ $(0.5\text{m})^3$ volume with 25 sources and 2200 detectors
- We study this system using thousands of “bags” generated by our stochastic bag generator
- Bags may be threats due to shapes and/or materials
- Measurement time (i.e., dwell time) is linearly related to photon number for current systems (especially significant for medical applications)

Non-Threat Bag Example

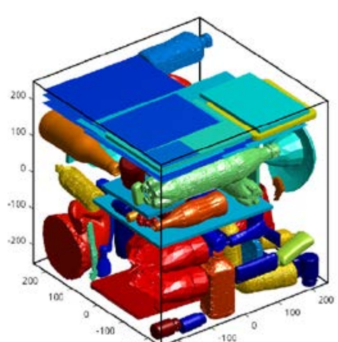

Threat Bag Example

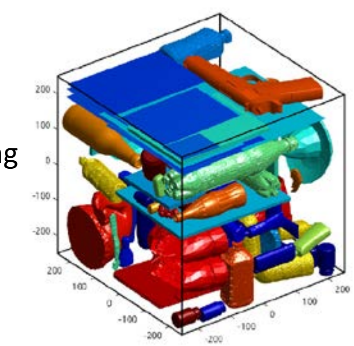

Compressive Reconstruction

Dictionary Constrained Nonlinear Reconstruction

- The basic idea here is that the sparse transformations are learned from the reconstructed image itself during reconstruction.

$$\hat{\mathbf{f}} = \arg\min_{\mathbf{f},\mathbf{D},\mathbf{A}} \|\mathbf{Hf} - \mathbf{g}\|_2 \ \text{ such that } \|\mathbf{DA} - \mathbf{Pf}\|_F^2 < \sigma, \|\alpha_i\|_0 < \tau$$

  D – learned dictionary
  A – collection of all dictionary coefficients
  α_i – dictionary coefficients of patch i
  P – reformats the object into a collection of patches (local spatial regions)
  σ, τ – dictionary error, and coefficient sparsity thresholds

- Solve by alternating nonlinear conjugate gradient and dictionary learning update:
  1. Update $\boldsymbol{f}$ from nonlinear conjugate gradient
  2. Update D_n and A_n
     a. $D_{n+1} = KSVD(D_n, \boldsymbol{f}_n, \sigma)$
     b. $A_{n+1} = OMP(D_{n+1}, \boldsymbol{f}_n, \tau)$
  3. Constrain $\boldsymbol{f}_n$ to its dictionary representation
     a. $\boldsymbol{f}_{n+1} = P^{-1}(D_{n+1}A_{n+1})$
  4. Iterate

Original

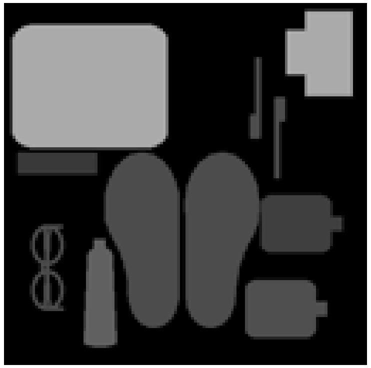

Low dimensional example requires ~ 150 views for perfect linear reconstruction
All results based on 10^9 photons

Conventional SART Reconstruction

30 Views

60 Views

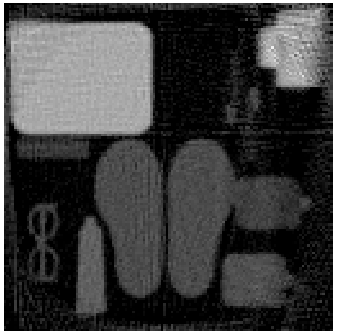

Dictionary Prior Reconstruction

30 Views

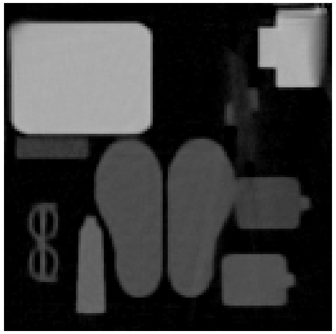

60 Views

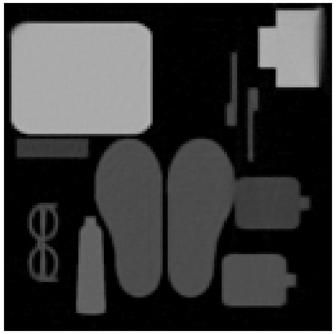

Threat Detection via Adaptive Tomography

- Multi-source x-ray architecture is ideally suited to adaptive measurement
- Static design exploits statistical knowledge of objects and task (design-time optimization).
- Adaptive design exploits additional knowledge obtained from previous measurements (on-line optimization).
- Greedy adaptation maximizes the benefit of the “next” measurement.
- Sequential hypothesis testing (SHT) is a established formalism for on-line Bayesian experimental design.

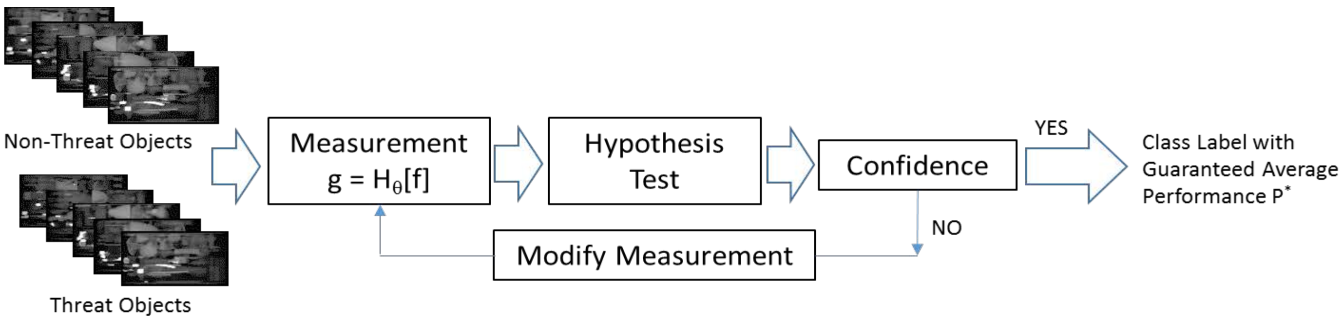

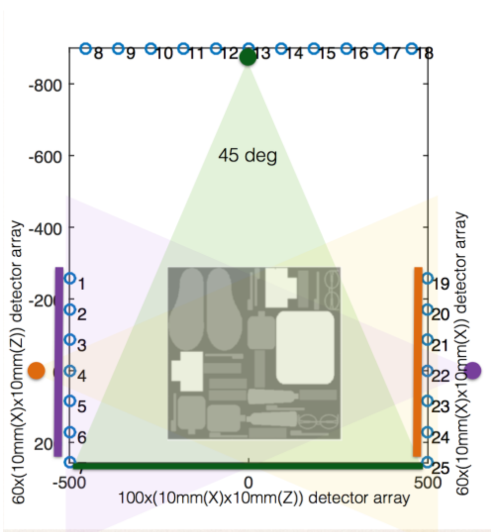

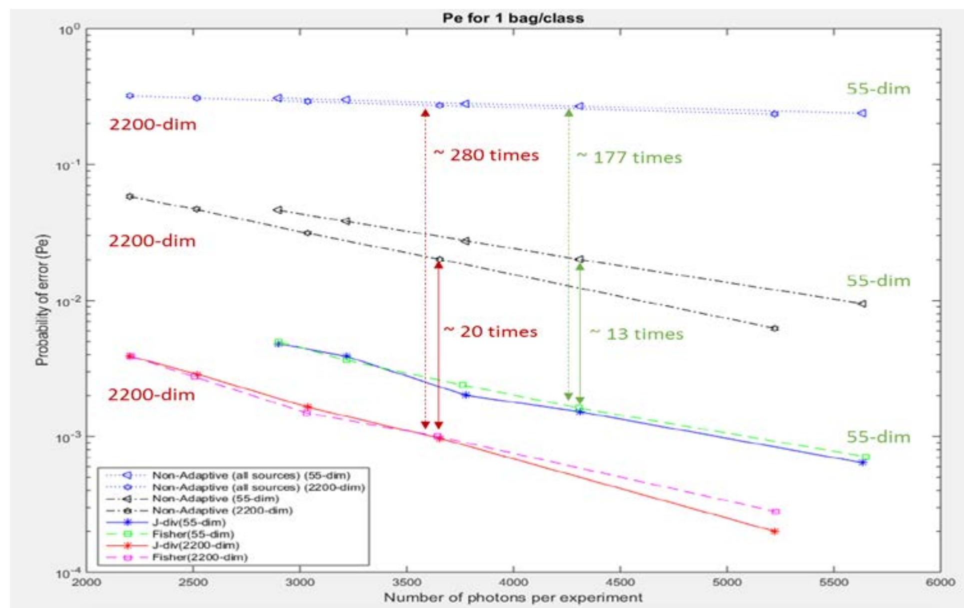

Conclusions

1. Computational imaging (i.e., joint design of hardware/algorithms) has emerged from the desire to leverage Moore's law toward optimizing sensor resources/costs.

2. Compressive measurement attempts to match the information content of a signal with the sensing resources required to measure that signal.
   a. Don't measure what you already know
   b. Don't measure what you don't need to know

3. This perspective is particularly important when
   a. Measurements are expensive
   b. Dimensionality mismatch between object and sensor

4. A task-specific definition of information is a useful tool for analysis/design and enables task-specific computational imagers to substantially outperform their conventional counterparts.

5. Some examples demonstrate the potential benefits of TSI-based design (VIS/IR, X-Ray, RF, …)