Integrated Imaging:

Creating Images from the Tight Integration of Algorithms, Computation, and Sensors

Charles A. Bouman
School of Electrical and Computer Engineering
Purdue University
CVPR Workshop on Computational Cameras and Displays (CCD)
2014-06-28

Co-authored with:

Ken Sauer, University of Notre Dame
Jean-Baptiste Thibault, GE Healthcare
Jiang Hsieh, GE Healthcare
Zhou Yu, GE Healthcare
Venkat Venkatakrishnan, Purdue
Larry Drummy, AFRL
Marc De Graef, CMU
Jeff Simmons, AFRL
Brendt Wohlberg, LANL
Peter Voorhees, Northwestern University
Greg Buzzard, Purdue University

- Supported by:
  - GE Healthcare
  - Air Force Office of Scientific Research, MURI contract # FA9550-12-1-0458, and the Air Force Research Laboratory

Integrated Imaging: Combining Algorithms and Physical Sensors

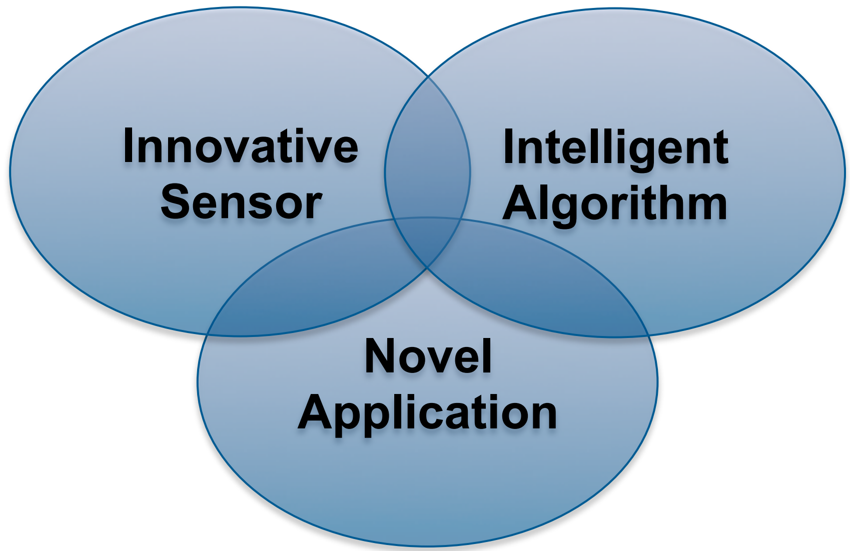

- Traditional sensor design is reaching its limits
  - Difficult to only measure one parameter
  - No longer possible to "fix" the device
- Rather than making the **"purest"** measurement, make the **most informative** measurement.
- Requires tight integration of sensor and algorithms.

Transactions on Computational Imaging (TCI)

- New IEEE journal
- Multidisciplineary systems oriented research
- Data in => images out
- Topics include:
  - Computational Photography
  - Computed imaging
  - Novel sensing systems
  - Applications in consumer imaging, biomedical imaging, remote sensing, scientific imaging, industrial imaging
- Will start early 2015
- We need your help and support!

Integrated Imaging: The Philosophy *

- *Mick Jagger's Theorem:*

  You can't always get what you want, but if you try sometimes, you might get what you need.

- What should you get (measure)?
  - Don't measure one thing at a time very precisely
  - Measure everything mixed together adaptively

- How do you form an image from what you get?
  - Use all available information to form the image
  - Combine measurements and prior knowledge

Inverse Problem: Example

- Forward model
  - Gravity
  - Fluid dynamics
  - Light propagation
  - Image formation
- Inversion
  - Illumination estimation
  - Shape from X
  - Inverse dynamics
  - Real world knowledge

- Inverse Solution: Something fell in the water

Model Based Iterative Reconstruction (MBIR): A General Framework for Solving Inverse Problems

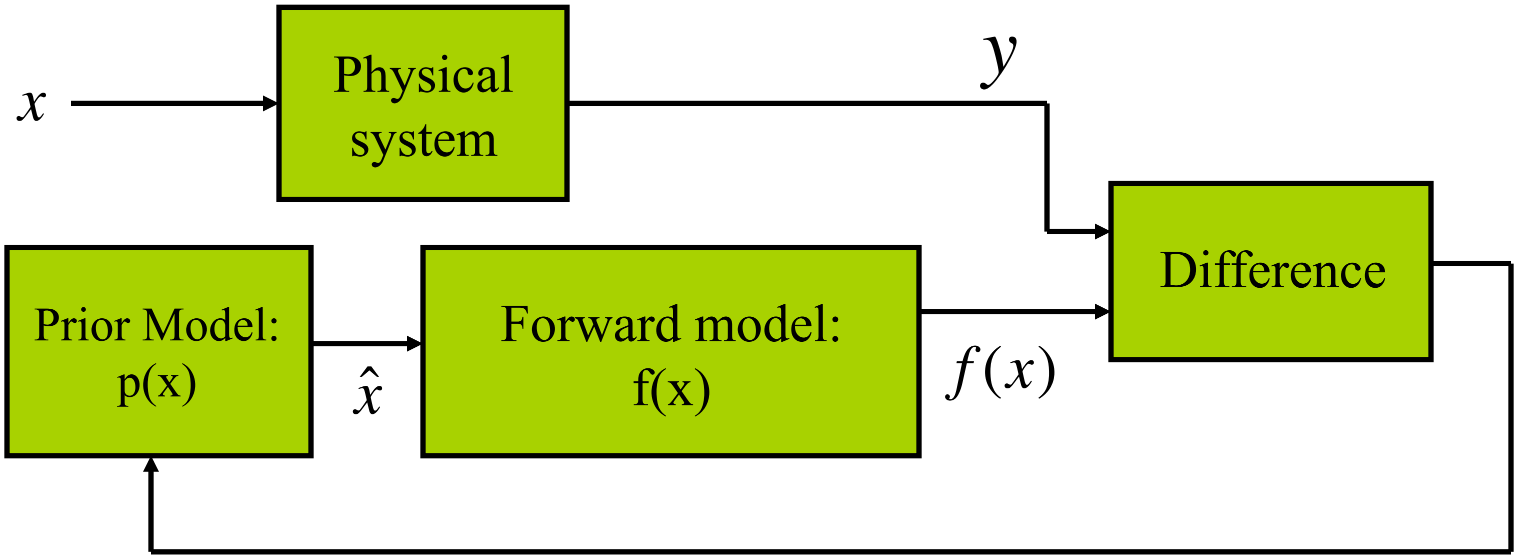

$$\hat{x} \leftarrow \arg\max_{x} \{ \underbrace{\log p(y \mid x)}_{\text{forward model}} + \underbrace{\log p(x)}_{\text{prior model}} \}$$

$\hat{x}$ – Reconstructed object

y – Measurements from physical system

"Thin Manifold" View of Prior Models

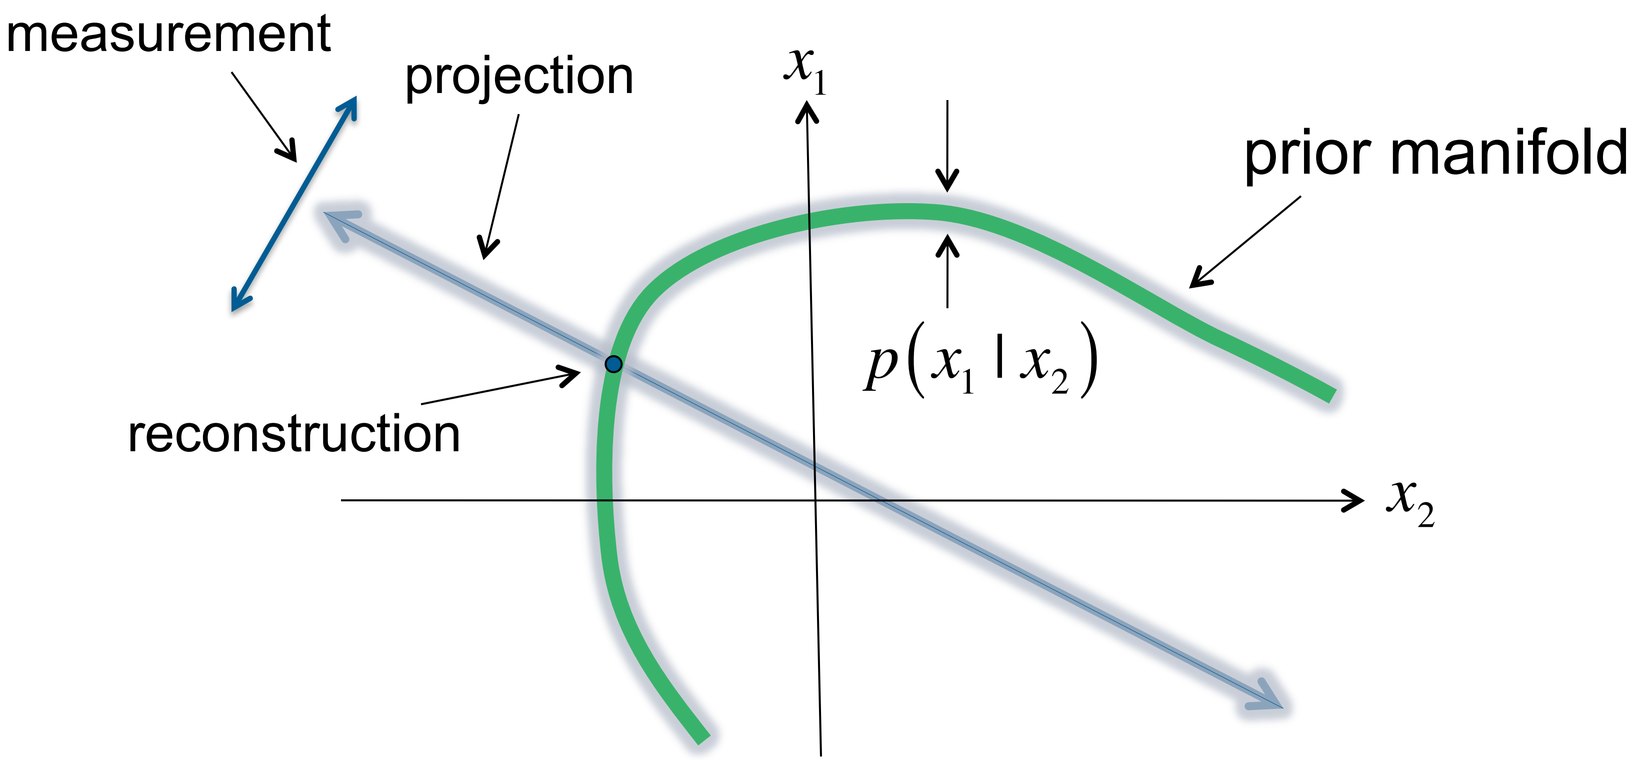

- Notice that prior manifold fills the space but…
  - Not a linear manifold
  - PCA can not effectively reduce dimension
- But it has thickness
- Dimension of measurement > dimension of manifold

Recipe for Integrated Imaging

- Design sensor to measure the most informative data
- Form image:
  - Solve inverse problem: Data in => image out
  - Synergy between **forward model of sensor** and **prior model of image**
- Explosion of possibilities:
  - Mix and match sensors and models
  - Do things you couldn’t do before

Medical CT Imaging

Ken Sauer, University of Notre Dame
Jean-Baptiste Thibault, GE Healthcare
Jiang Hsieh, GE Healthcare

Multi-slice Helical Scan Medical CT

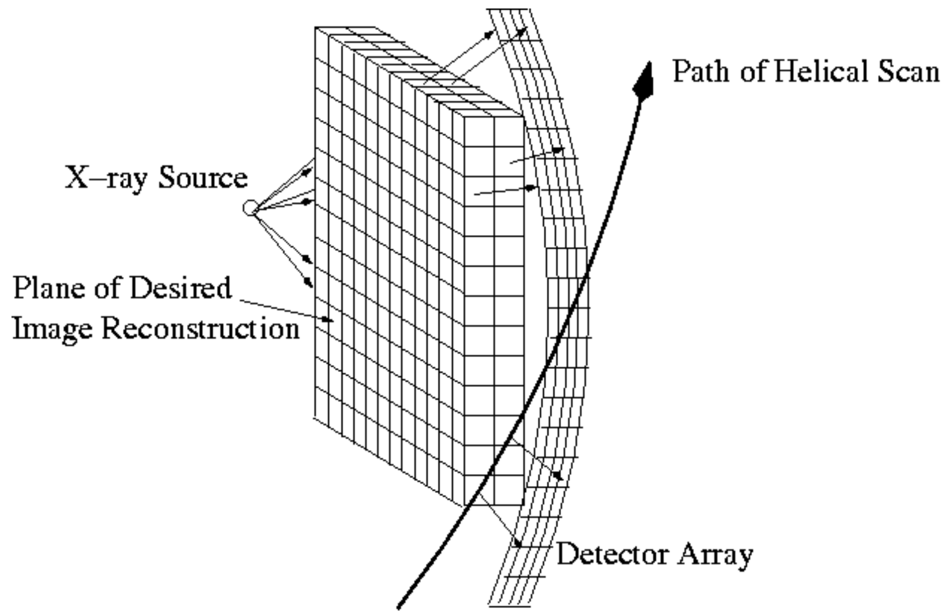

- Reconstruct 3D volume form 1D projections
- Geometry:
  - Helical scan
  - Multislice => cone angle in 3D

Model-Based Iterative Reconstruction (MBIR)

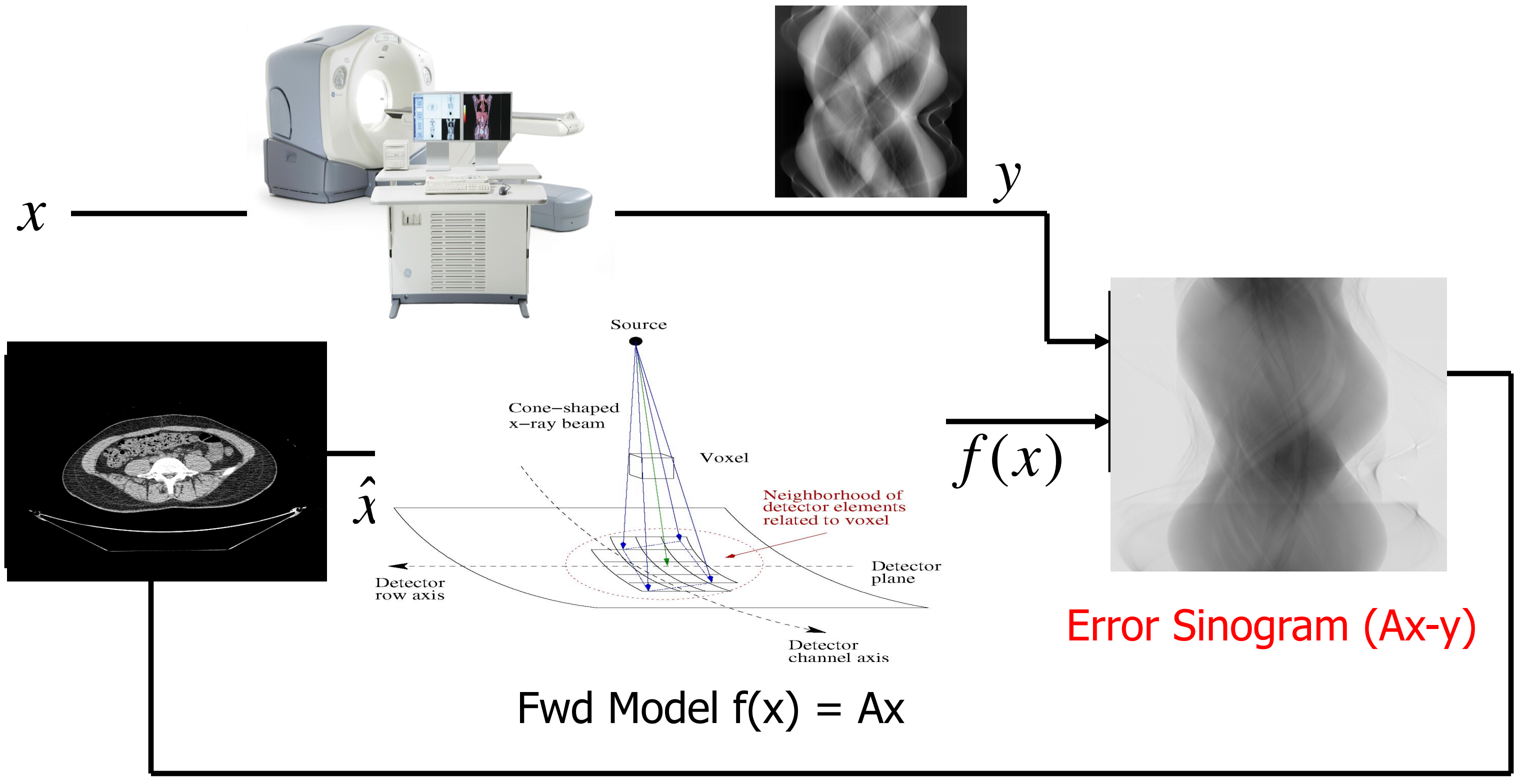

$$\hat{x} = \arg\min_{x\geq 0}\left\{-\log p\left(y \mid x\right) - \log p\left(x\right)\right\}$$

$$= \arg\min_{x\geq 0}\left\{\frac{1}{2}\left\|y - Ax\right\|_{\Lambda}^{2} \quad + \quad u\left(x\right)\right\}$$

CT Scanner Forward Model: *p(y|x)*

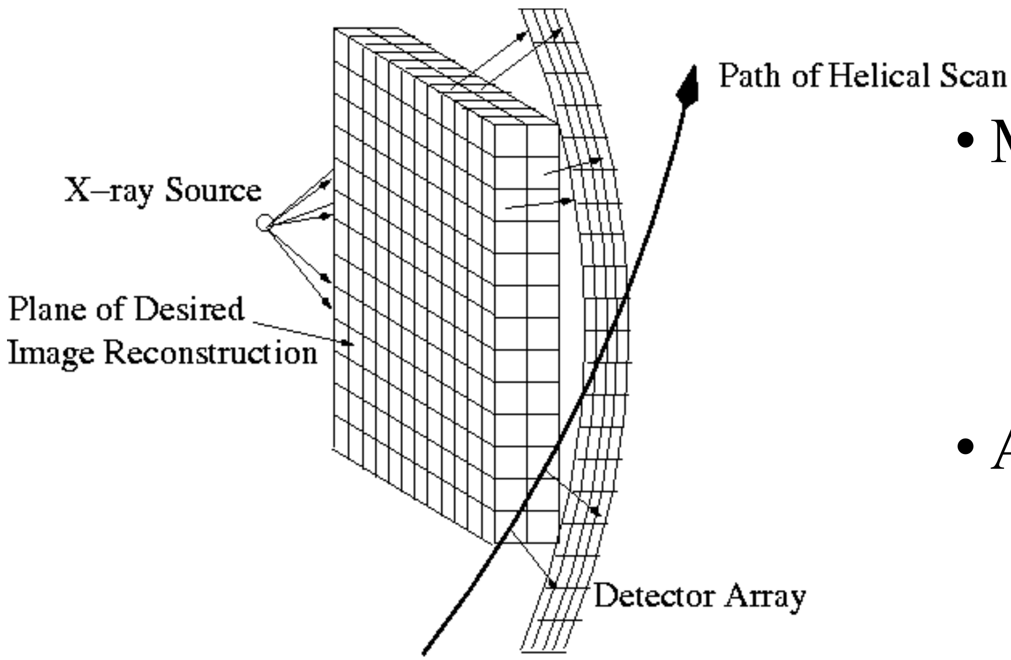

- Measure X-ray attenuation

$$y_i = -\ln\left(\frac{\lambda_i}{\lambda_T}\right)$$

- Assumptions

$$E\left[y \mid x\right] = Ax$$

$$Var\left[y \mid x\right] = \Lambda$$

Important! $\Rightarrow$ $Var\left[y_i \mid x\right] = \Lambda_{i,i} = \frac{\lambda_i + \sigma_e^2}{\lambda_i^2}$ – Photon counting + electronic noise

Forward model:

$$-\log p(y \mid x) = \frac{1}{2}\left\|y - Ax\right\|_\Lambda^2 + \text{ constant}$$

Markov Random Field (MRF) Prior Model

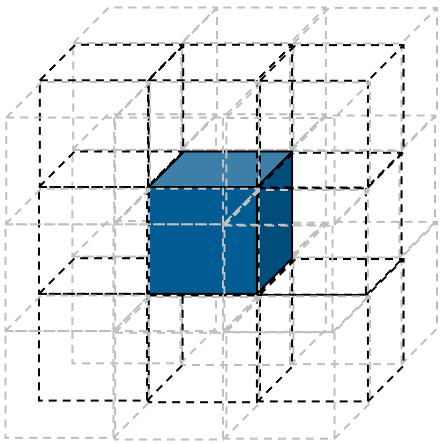

3D Neighbors

- Gibbs Distribution

$$p(x) = \frac{1}{Z} \exp\left\{ -\sum_{\{j,k\} \in C} \rho\left(\frac{x_j - x_k}{\sigma} \right) \right\}$$

$\rho\left(x_j - x_k\right)$: Potential function

- Properties:
  - MRF with 26 local neighbors in 3D
  - $\rho(\Delta)$ preserves edges

Prior model:

$$-\log p(y \mid x) = -\sum_{\{j,k\} \in C} \rho\left(\frac{x_j - x_k}{\sigma} \right)$$

Choice of MRF Potential Function

$\rho(f_i - f_j)$: Penalty on the difference between neighboring voxels

$$\text{If } \rho(f_i - f_j) = \frac{\left|\frac{f_i - f_j}{\sigma_f}\right|^2}{c + \left|\frac{f_i - f_j}{\sigma_f}\right|^{2-p}}$$

$q -$ Generalized Gaussian MRF*

$p = 2$ corresponds to diffuse interfaces

$p = 1$ corresponds to sharp interfaces
- Total Variation Regularization
(compressed sensing)

σ_f : MRF scaling parameter (controls noise)

$\rho(f_i - f_j)$

$p = 2$

$p = 1$

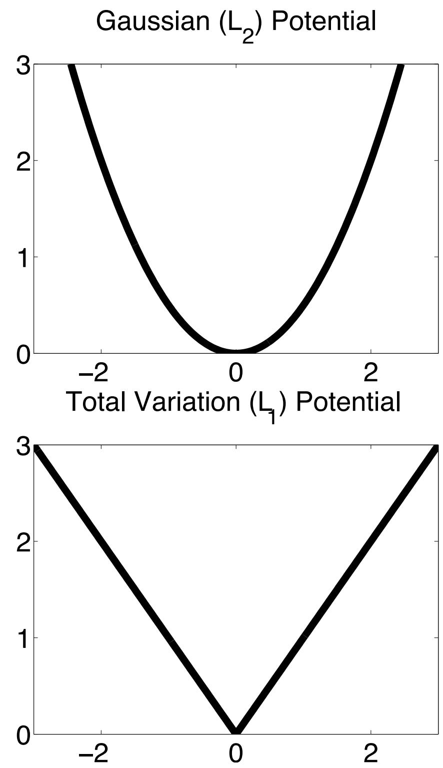

*J.-B. Thibault, K. Sauer, C. Bouman, and J. Hsieh, "A three-dimensional statistical approach to improved image quality for multi-slice helical CT," Med. Phys., vol. 34, no. 11, pp. 4526–4544, 2007

Optimization for MAP Estimation

$$\hat{x} = \arg\min_{x \geq 0} \left\{ \frac{1}{2}(y - Ax)^T \Lambda (y - Ax) + u(x) \right\}$$

- Approaches
  - ICD – fast robust convergence, but not so GPU friendly
  - Gradient based optimization – GPU friendly, but more fragile
- Other important tricks:
  - Non-homogeneous updates
  - Preconditioning
  - Ordered subsets
  - Nested optimization
  - Multiresolution/Targetting

RMSE Convergence Plots for NH-ICD

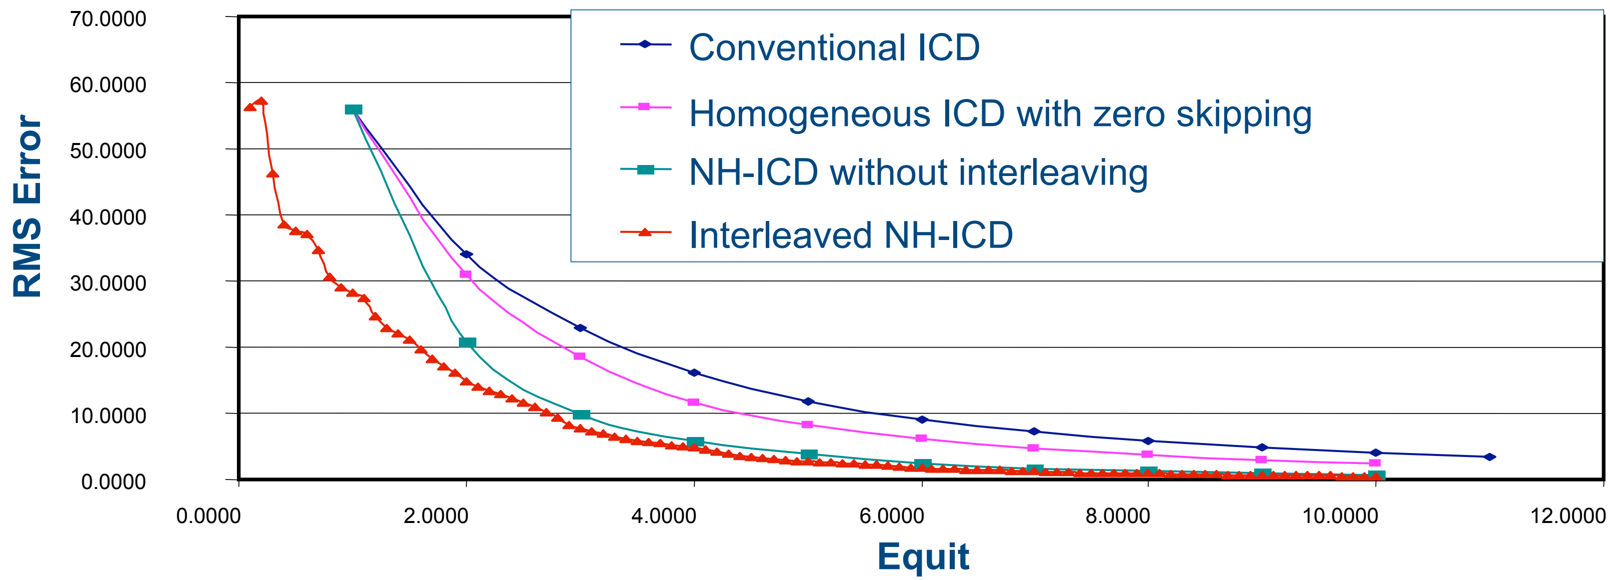

- NH-ICD
  - Reduces transients at early stage allowing faster convergence
  - Interleaving in early iterations further improves convergence speed

Zhou Yu, Jean-Baptiste Thibault, Charles A. Bouman, Ken D. Sauer, and Jiang Hsieh, "Fast Model-Based X-ray CT Reconstruction Using Spatially Non-Homogeneous ICD Optimization," to appear in the *IEEE Trans. on Image Processing*.

Model-Based Iterative Reconstruction (MBIR): GE Healthcare's Veo System

- What is Veo?
  - GE announce new product, "Veo", based on MBIR reconstruction at RSNA 2010
  - System received FDA 510(k) approval in 2011
  - Currently on sale in US as an upgrade option
  - Partnership between GE Healthcare, Purdue University and the University of Notre Dame
  - Research team:
    - Jean-Baptist Thibault, Jiang Hsieh (GE)
    - Ken Sauer (Notre Dame)
    - Me (Purdue)

Resolution vs Noise

GEPP wire, 16x0.625mm, P15/16:1, 100mA, 10cm fov

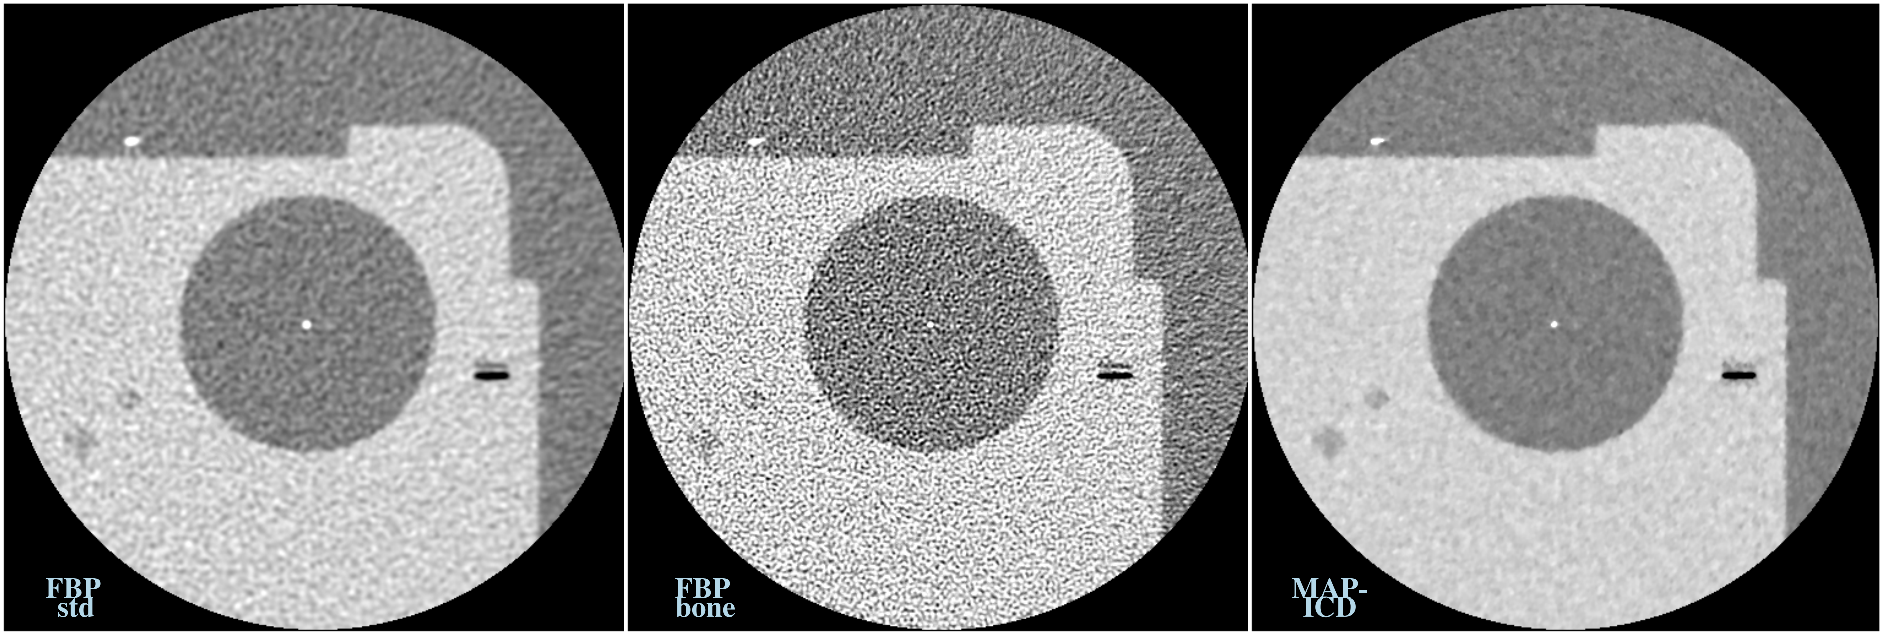

MTF comparable to FBP bone
50% lower noise than FBP std
Challenges usual trade-off

IQ	FBP std	FBP bone	MAP-ICD
50% MTF	4.39	8.53	8.66
10% MTF	7.04	11.90	13.20
Std dev	24.99	90.94	13.01

MBIR for 64 slice GE VCT Data

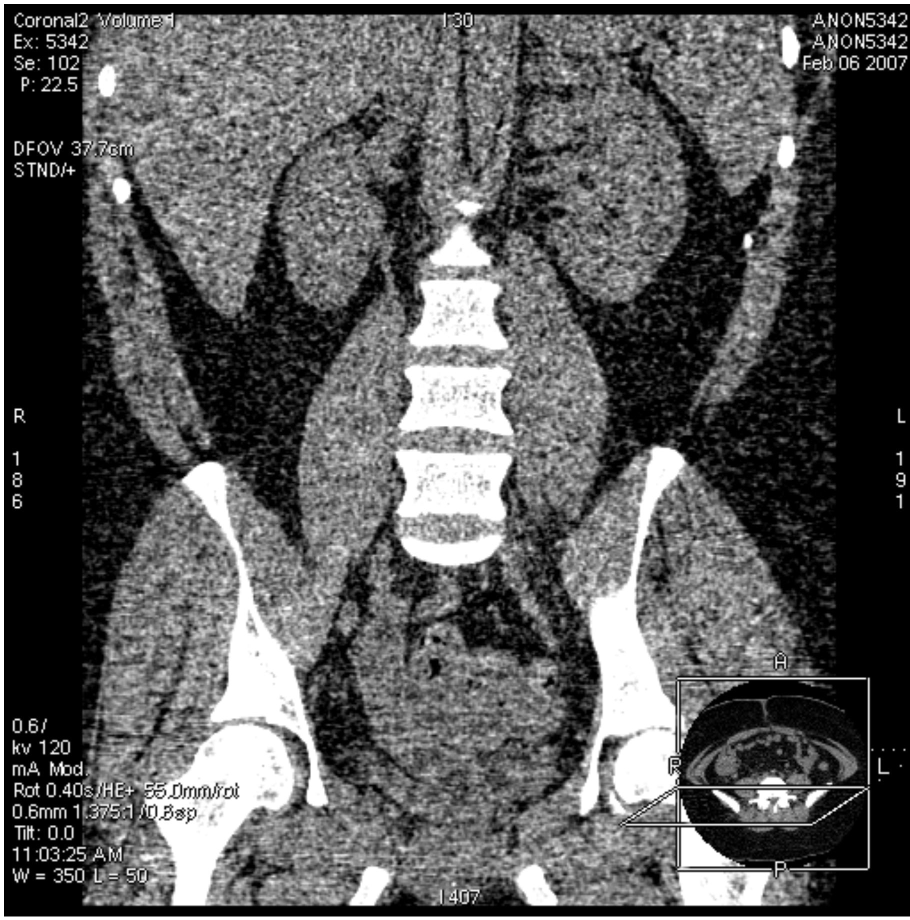

State-of-the-art 3D Recon

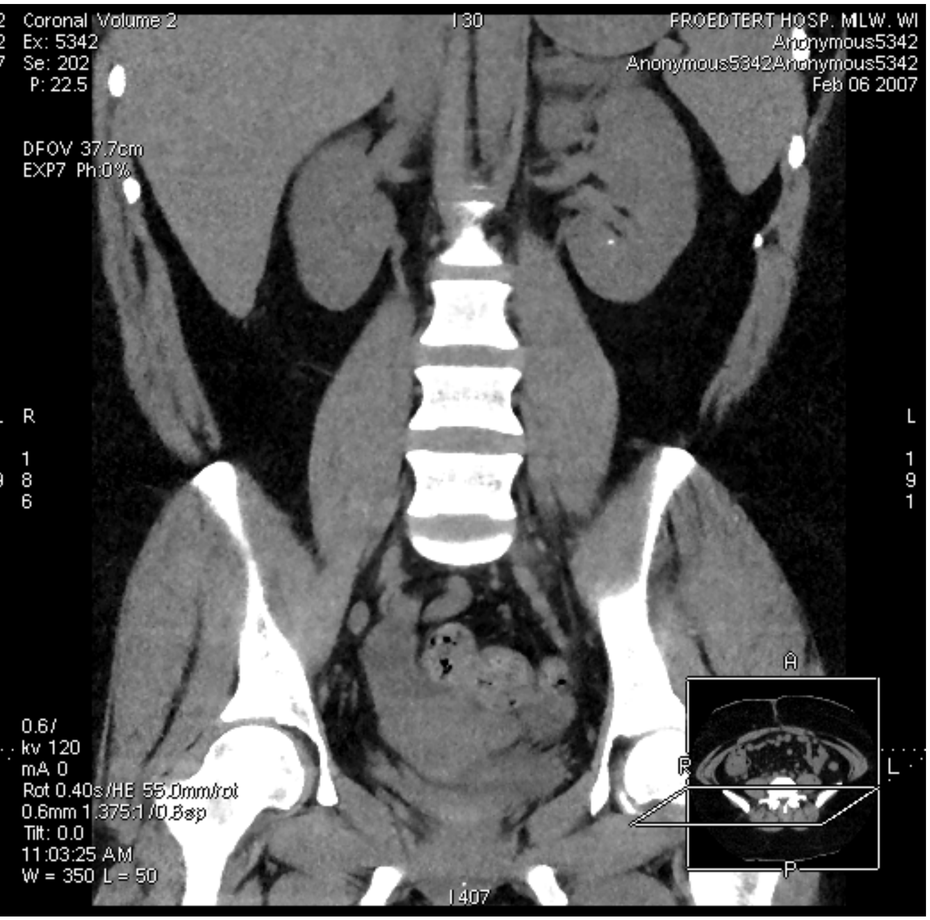

GE MBIR
Purdue/Notre Dame/GE algorithm

MBIR for 64 slice GE VCT Data

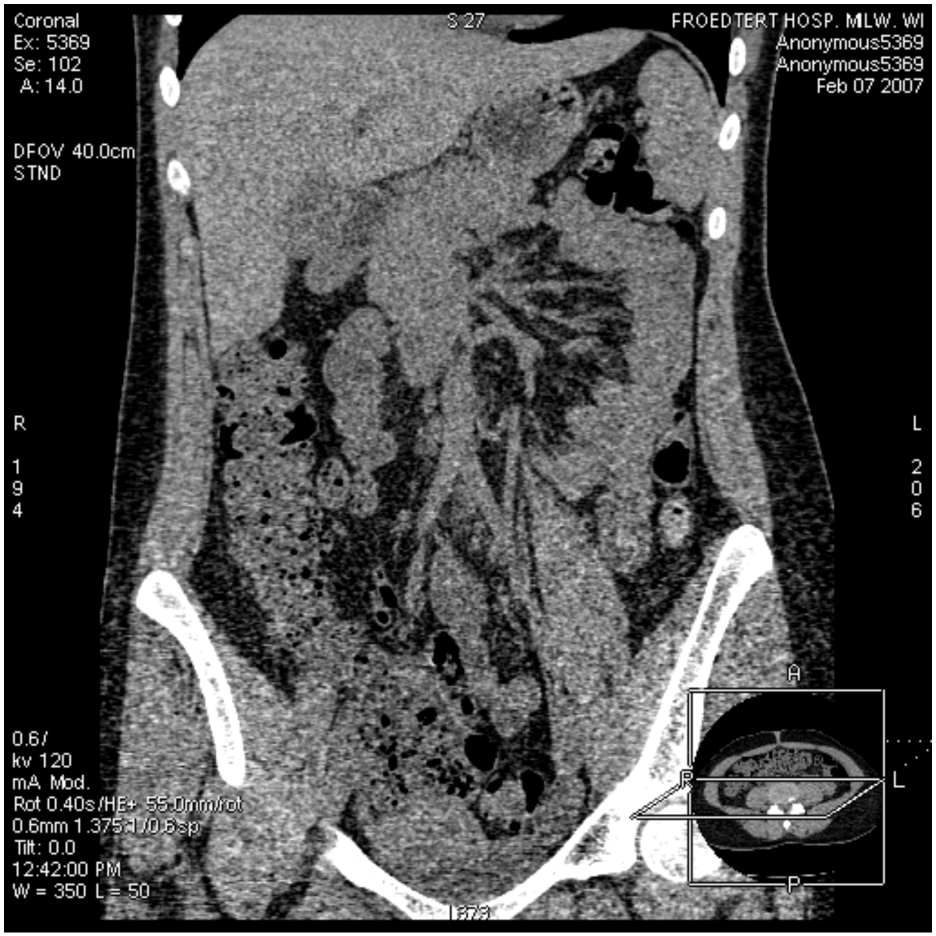

State-of-the-art 3D Recon

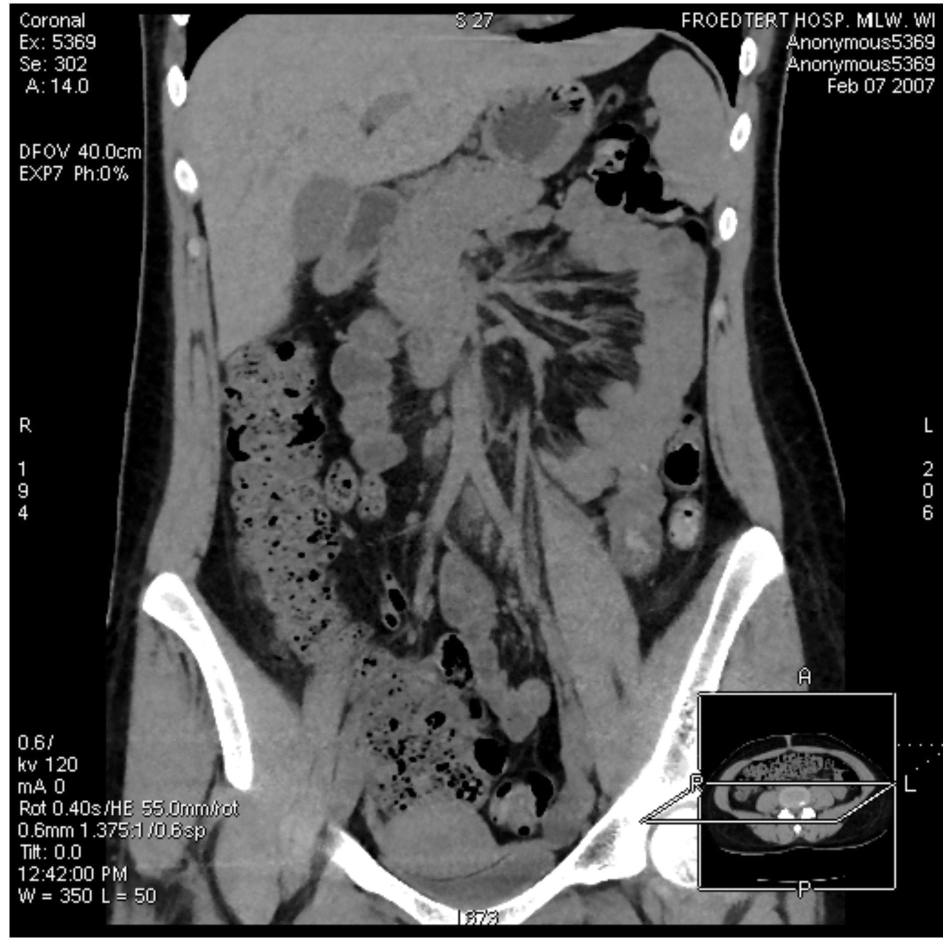

GE MBIR
Purdue/Notre Dame/GE algorithm

MBIR for 64 slice GE VCT Data

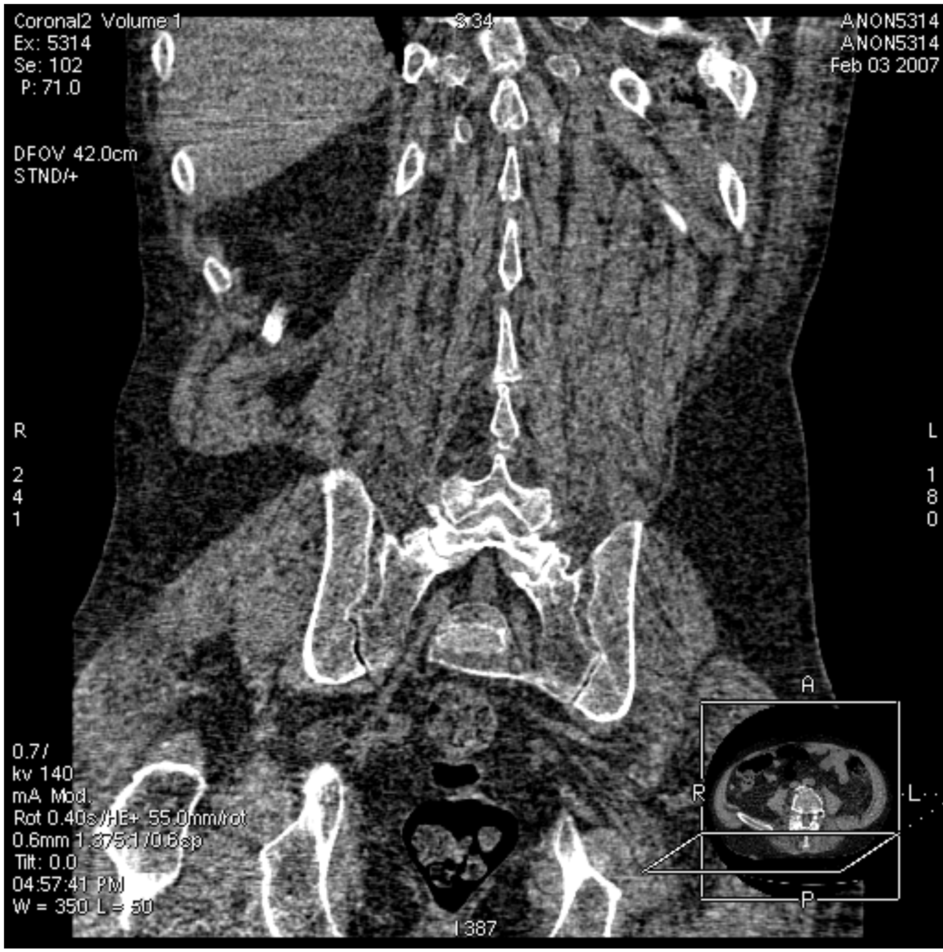

State-of-the-art 3D Recon

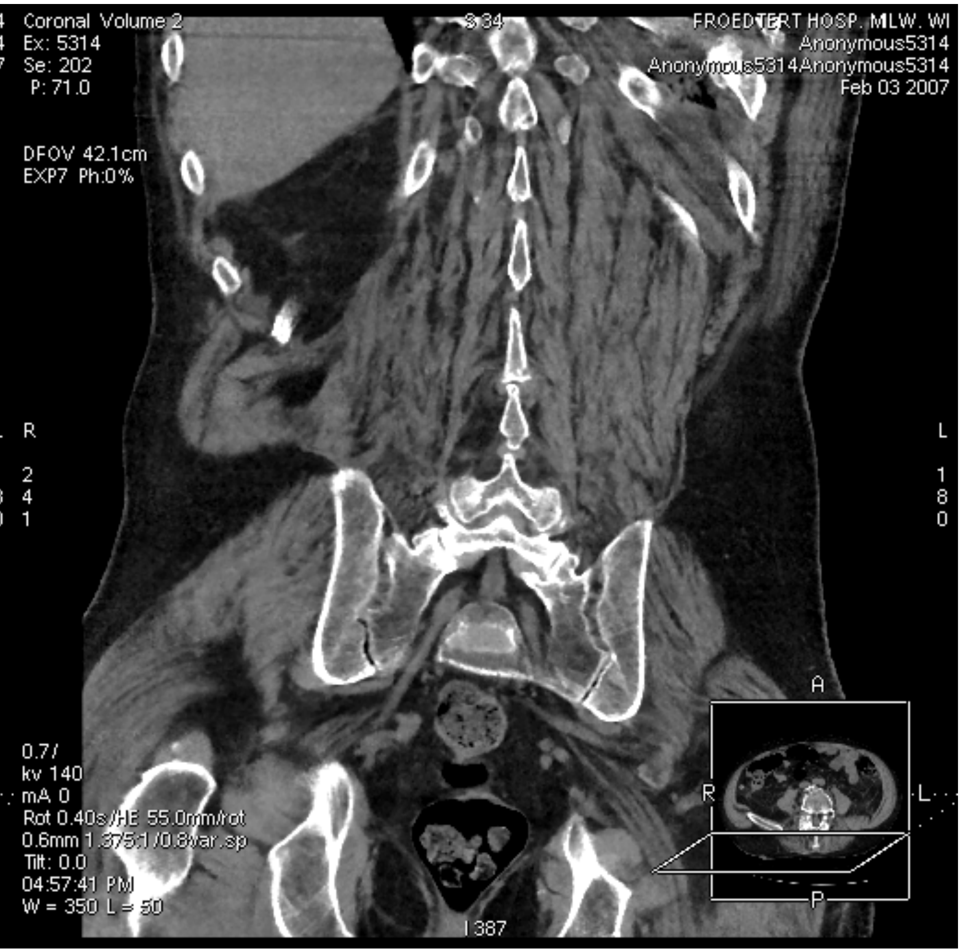

GE MBIR
Purdue/Notre Dame/GE algorithm

Pediatric Image at Low Dose (Coronal)

ASiR Reconstruction

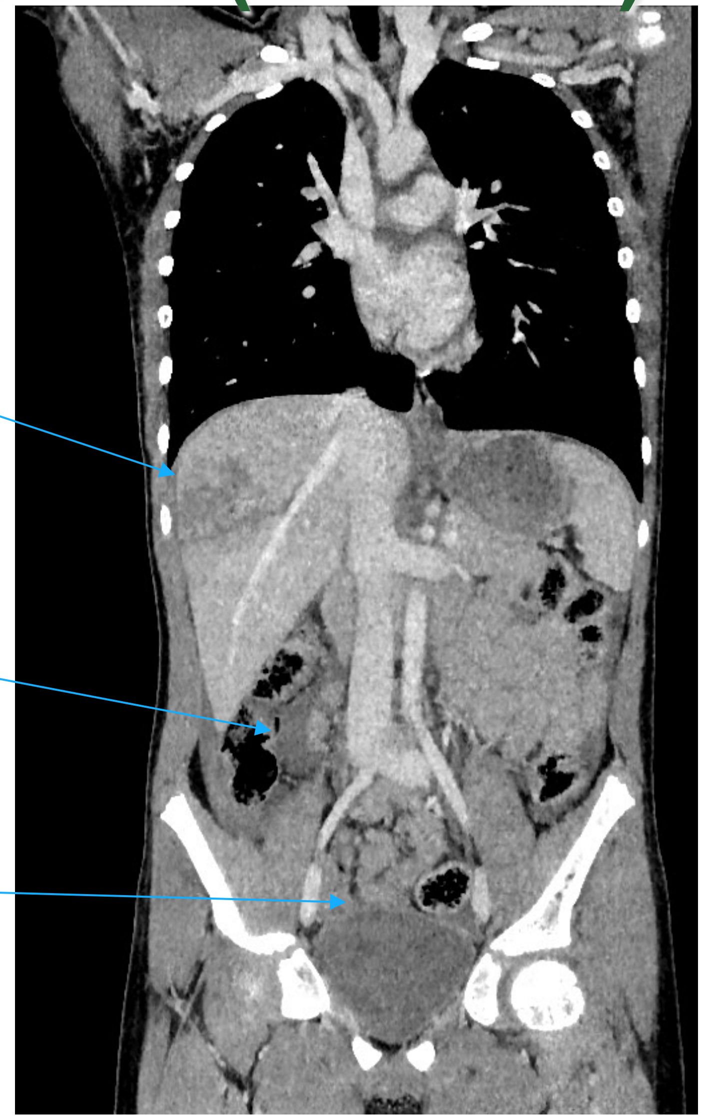

MBIR Reconstruction

Images courtesy of The Queen Silvia Children's Hospital

Dr. Stålhammar

VÄSTRA GÖTALANDSREGIONEN

Pediatric trauma, 120kV, 52-70mA, 0.4s/rot, 0.625mm, WW 300 WL 50

Pediatric Image at Low Dose (Transverse)

ASiR Reconstruction

MBIR Reconstruction

Images courtesy of The Queen Silvia Children's Hospital
Dr. Stålhammar

VÄSTRA GÖTALANDSREGIONEN

Pediatric trauma, 120kV, 52-70mA, 0.4s/rot, 0.625mm, WW 300 WL 50

Abdomen Imaging

Adrenal nodule

FBP Reconstruction

MBIR Reconstruction

kV 120, mA 150, 0.5s, 0.625mm, WW 350 WL 50 DFOV 42 Standard kernel in FBP

Images courtesy of Dr Gladys Lo

養和醫院 Hong Kong Sanatorium & Hospital

Time Interlaced Model Based Iterative Reconstruction (TIMBIR)

K. Aditya Mohan, Purdue
John Gibbs, NW
Prof. Peter Voorhees, NW
Prof. Marc De Graef, CMU
Dr. Xianghui Xiao, APS
Prof. Charles Bouman, Purdue

Synchrotron Imaging

An aerial view of the APS synchrotron at Argonne National Laboratory (Argonne, IL).

- Why are they important?
  - Intense, columnated, monochromatic source of X-rays
  - Have become more widely available
- Facilities
  - Advanced Photon Source (APS), Argonne National Labs; Advance Light Source (ALS), Lawrence Berkeley Labs; Cornell High Energy Synchrotron Source (CHESS); Stanford Synchrotron Radiation lightsource (SLAC); National Synchrotron Light Source, Brookhaven; Swiss Light Source.

Synchrotron Imaging of Time-Varying Sample

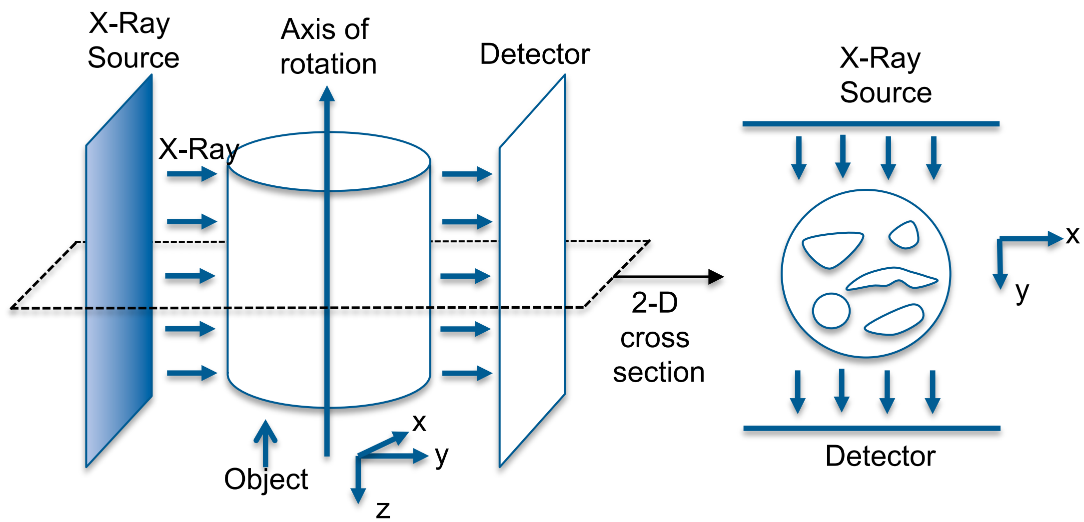

Real Synchrotron Projection Data

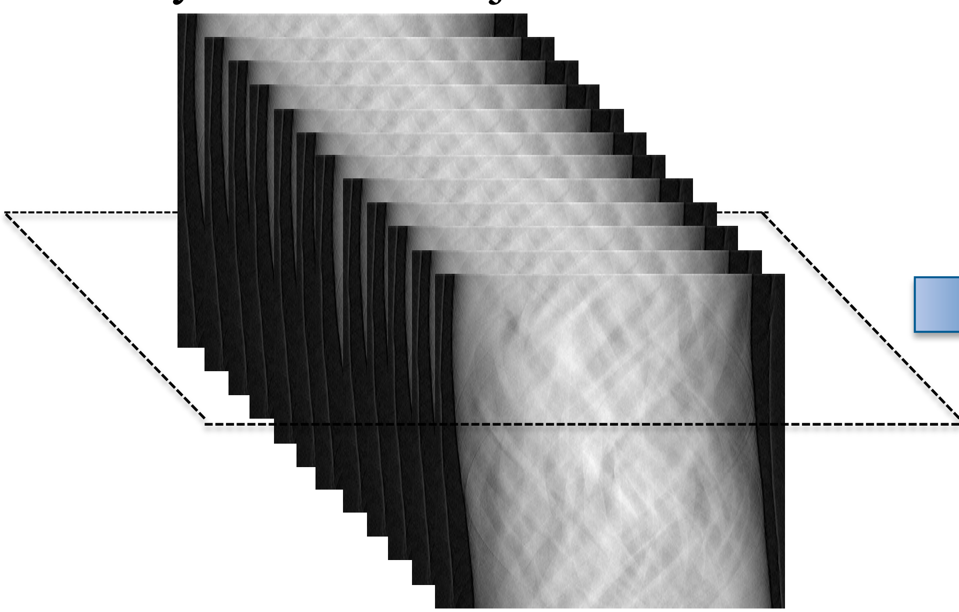

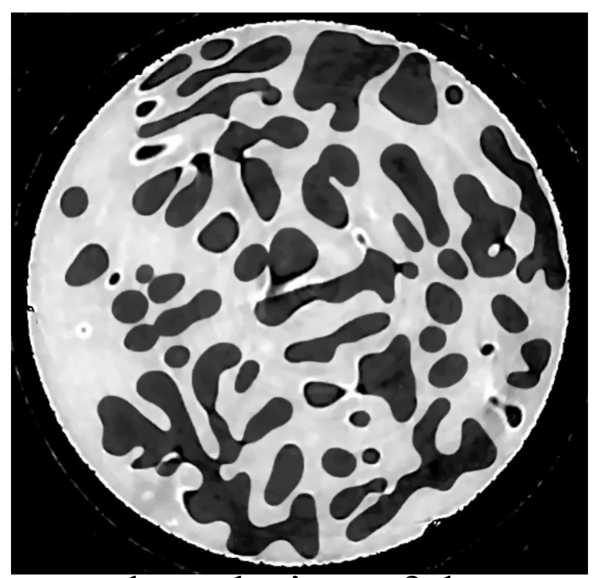

Temporal evolution of the sample

Conventional Approach to 4D Synchrotron Imaging

- Traditional approach
  - Acquire $N_v=2000$ views; do FBP reconstruction; repeat
  - Reduces time resolution by $N_v=2000$!!
- How do we increase temporal resolution ?

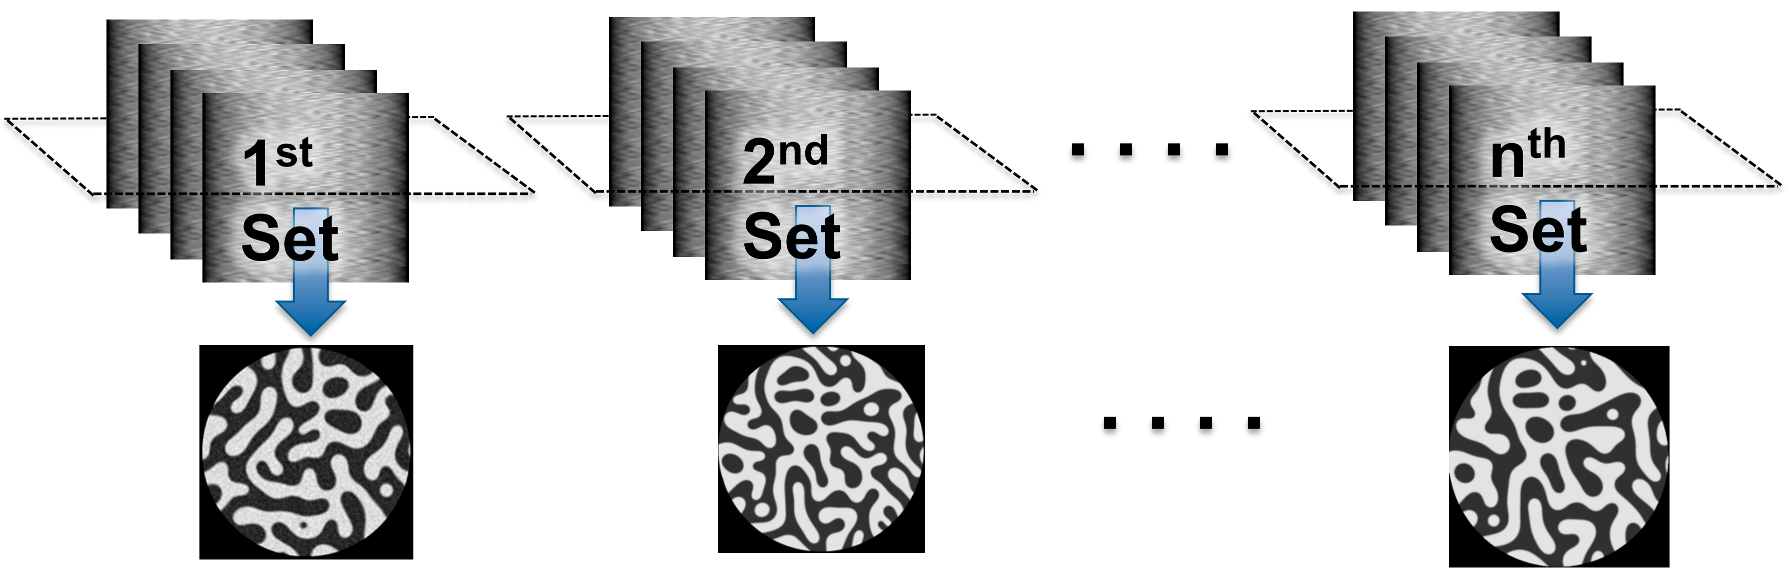

TIMBIR: Time Interlaced Model Based Iterative Reconstruction

- Interlace the views over K rotations of the object.

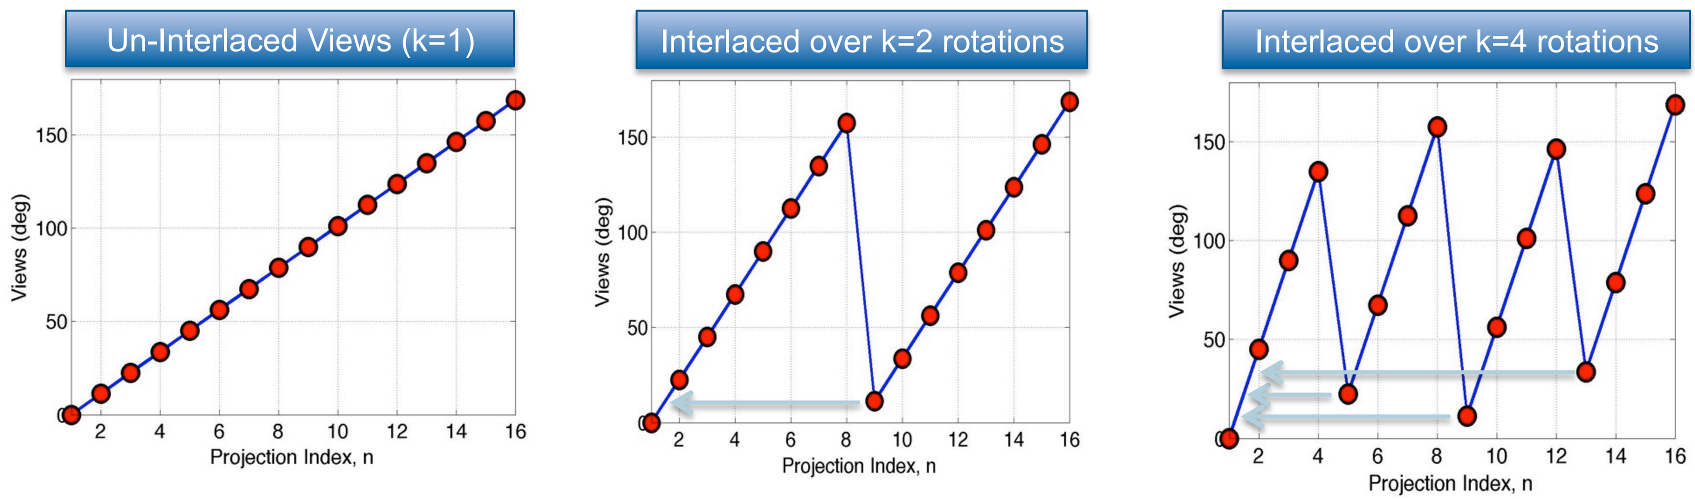

- Perform 4D MBIR reconstruction at any desired temporal resolution.

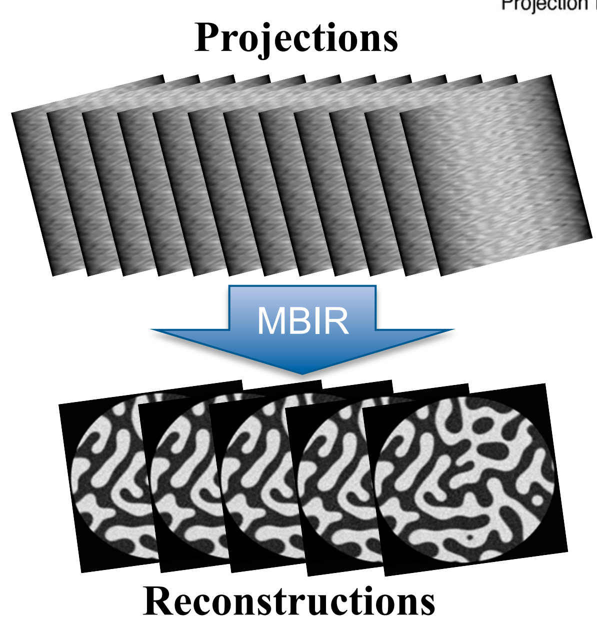

Examples of Interlaced Views

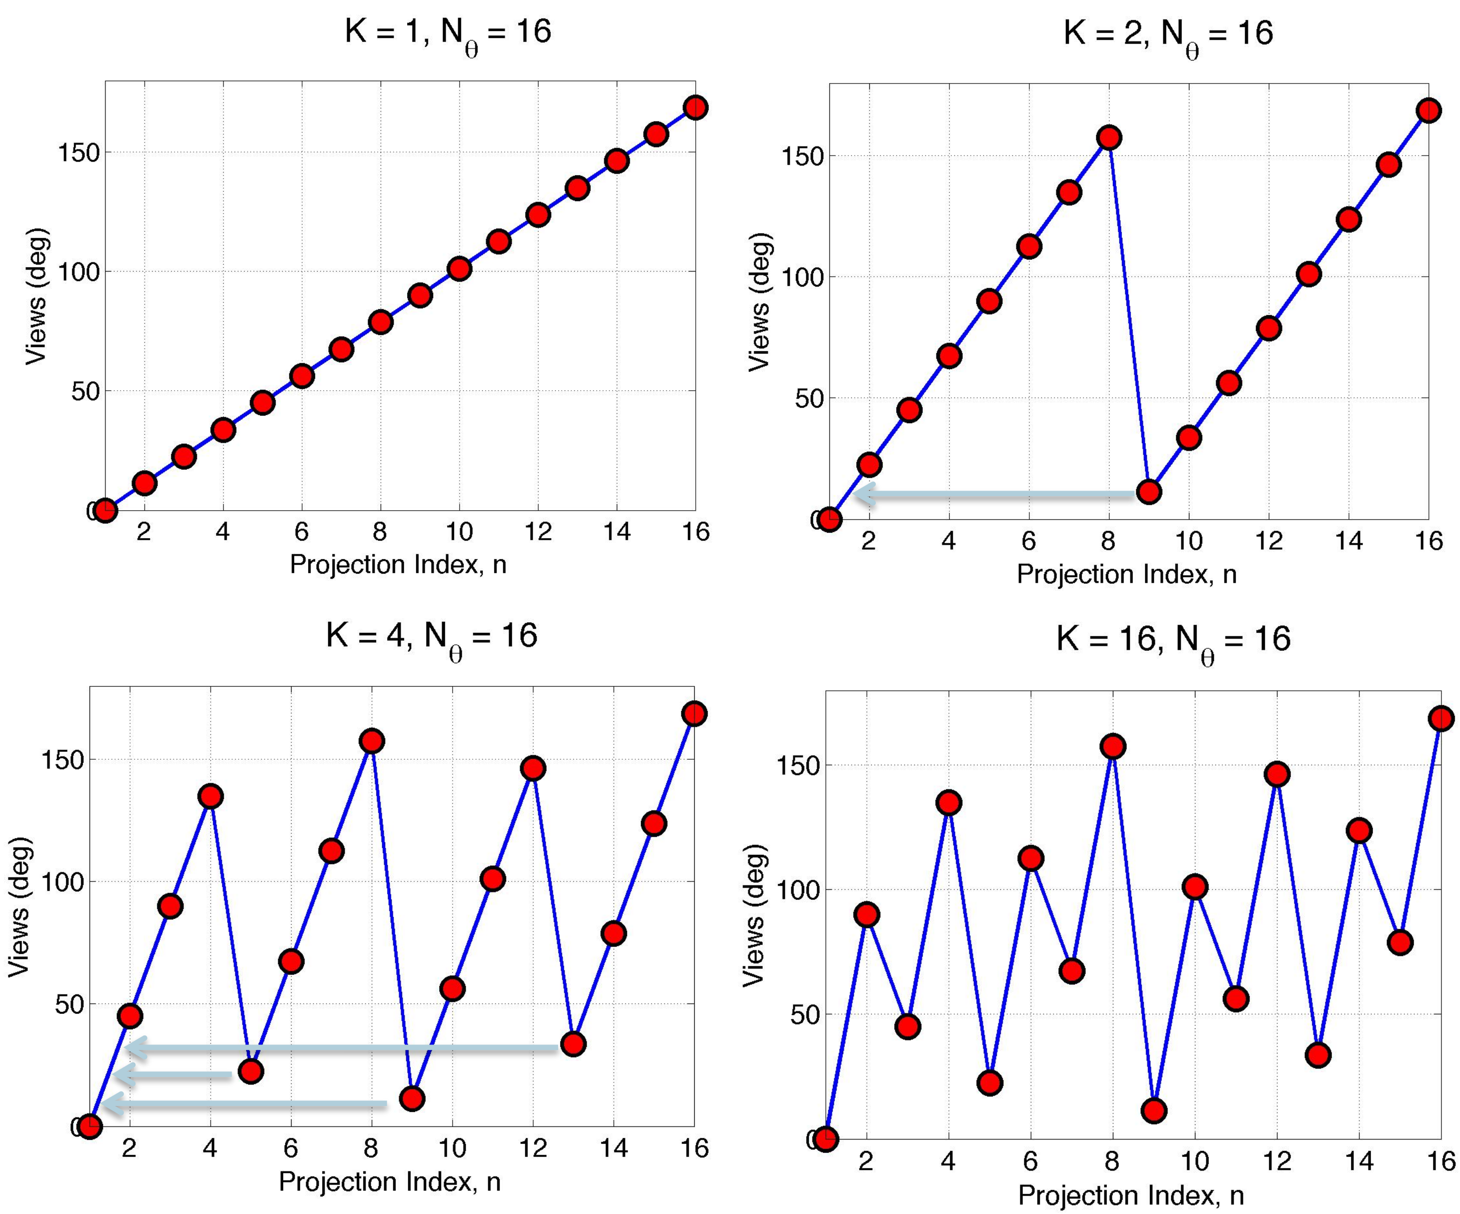

- Total number of discrete angles used is a constant.
- The time taken for rotation of object by 180 degrees decreases as K increases (or L decreases).

Number of Reconstructions per Frame, *r*

Case 1

$r = 1$

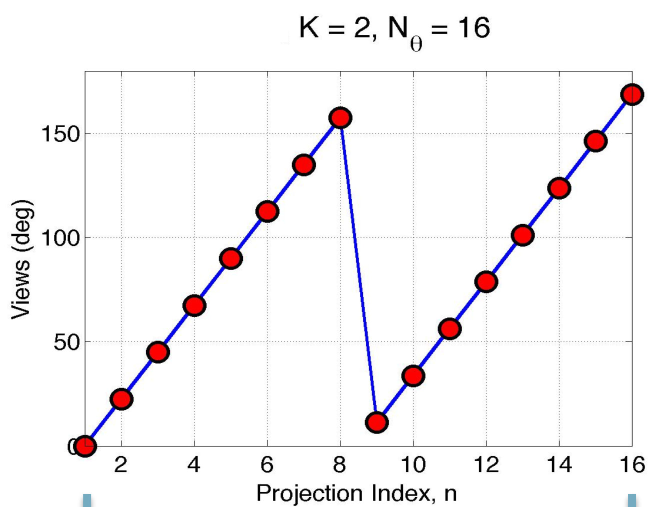

1 reconstruction for every 16 projections.

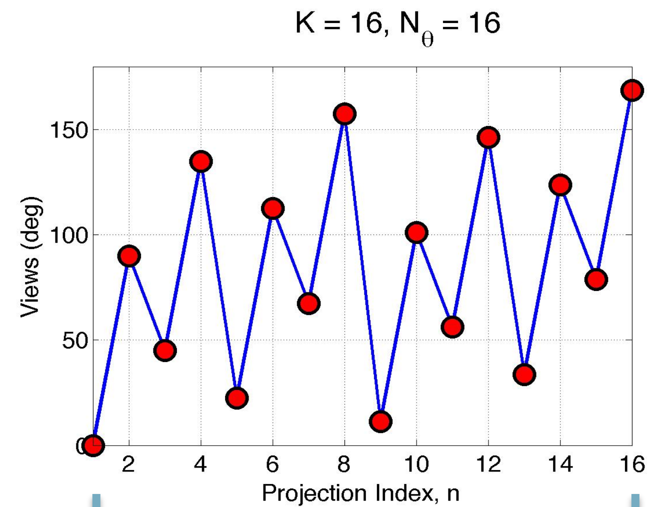

1 reconstruction for every 16 projections.

Case 2

$r = 2$

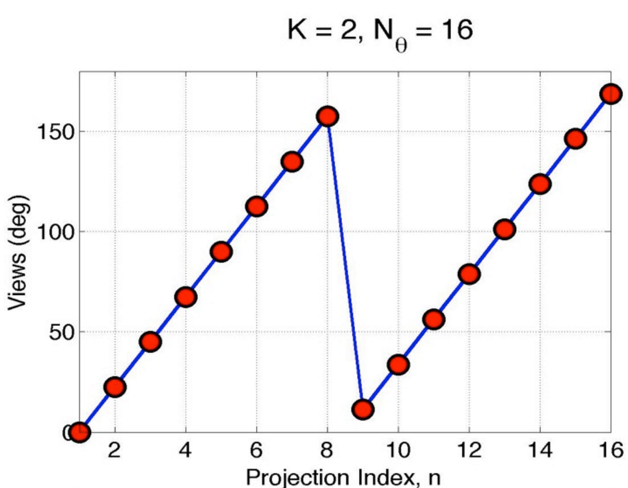

1 reconstruction for every 8 projections.

1 reconstruction for every 8 projections.

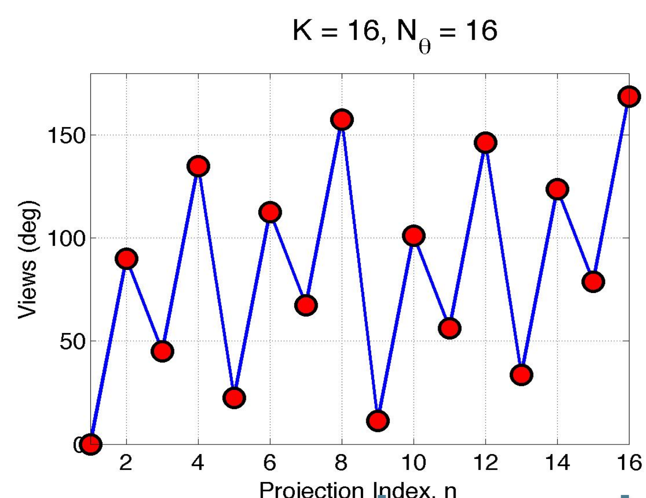

1 reconstruction for every 8 projections.

1 reconstruction for every 8 projections.

3D Reconstructions – FBP vs. TIMBIR

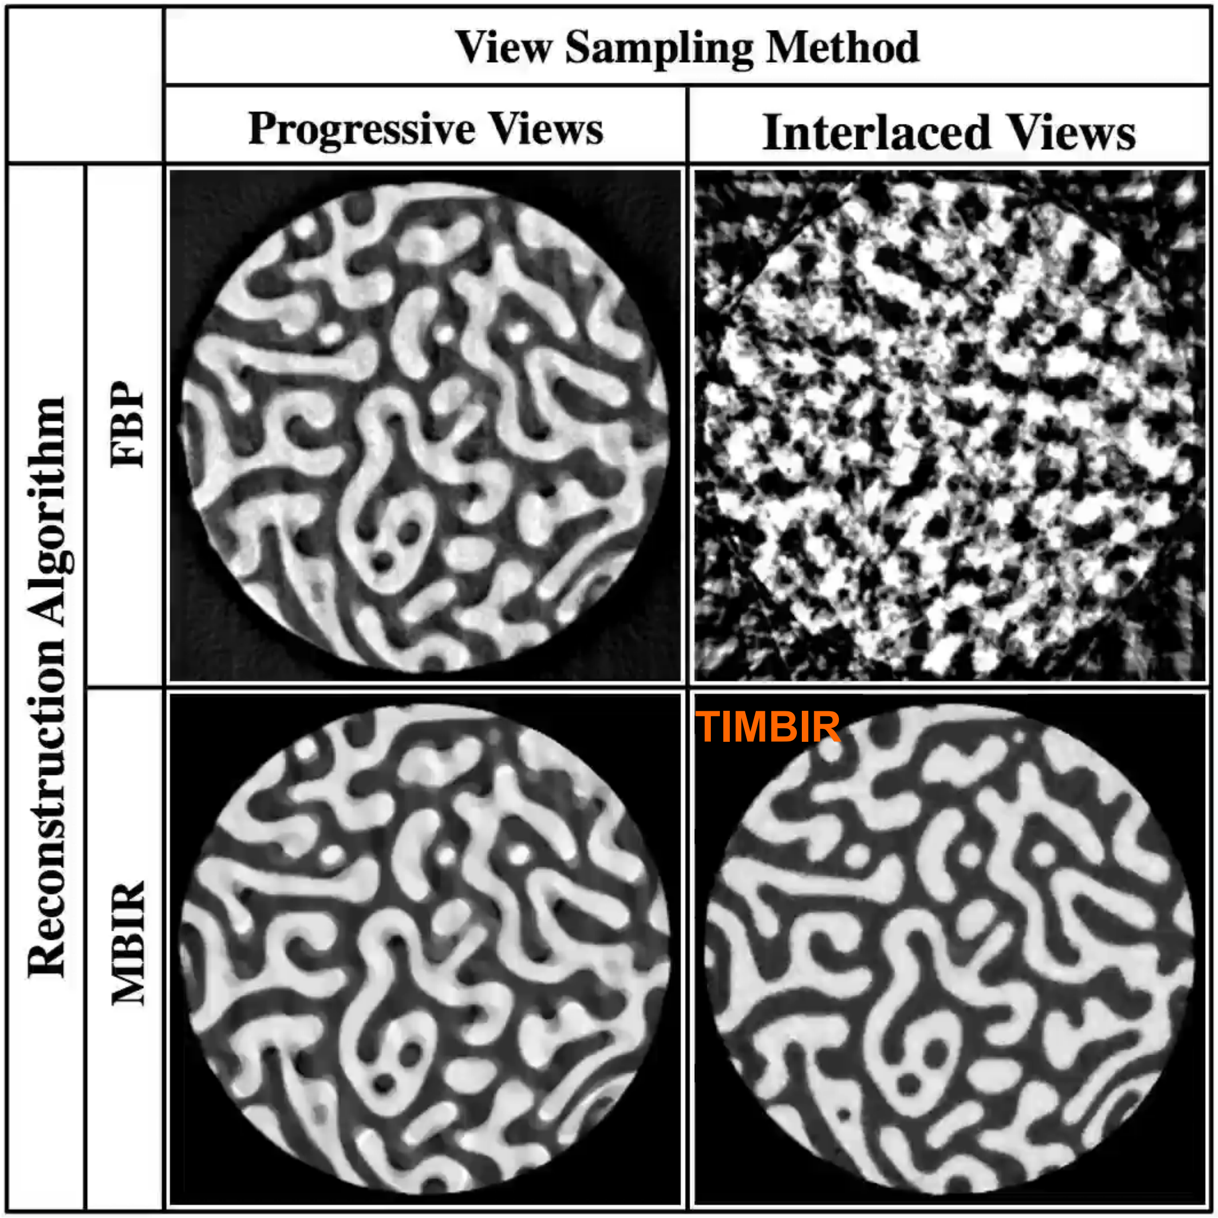

Method	RMSE (μm^{-1})
FBP/Progressive r = 1, K = 1, N_θ = 256	0.2528
FBP/Interlaced r = 1, K = 1, N_θ = 256	0.5867
MBIR/Progressive r = 1, K = 1, N_θ = 256	0.1032
MBIR/Interlaced r = 16, K = 16, N_θ = 256	0.0853

TIMBIR vs. FBP- Y Axis Slice

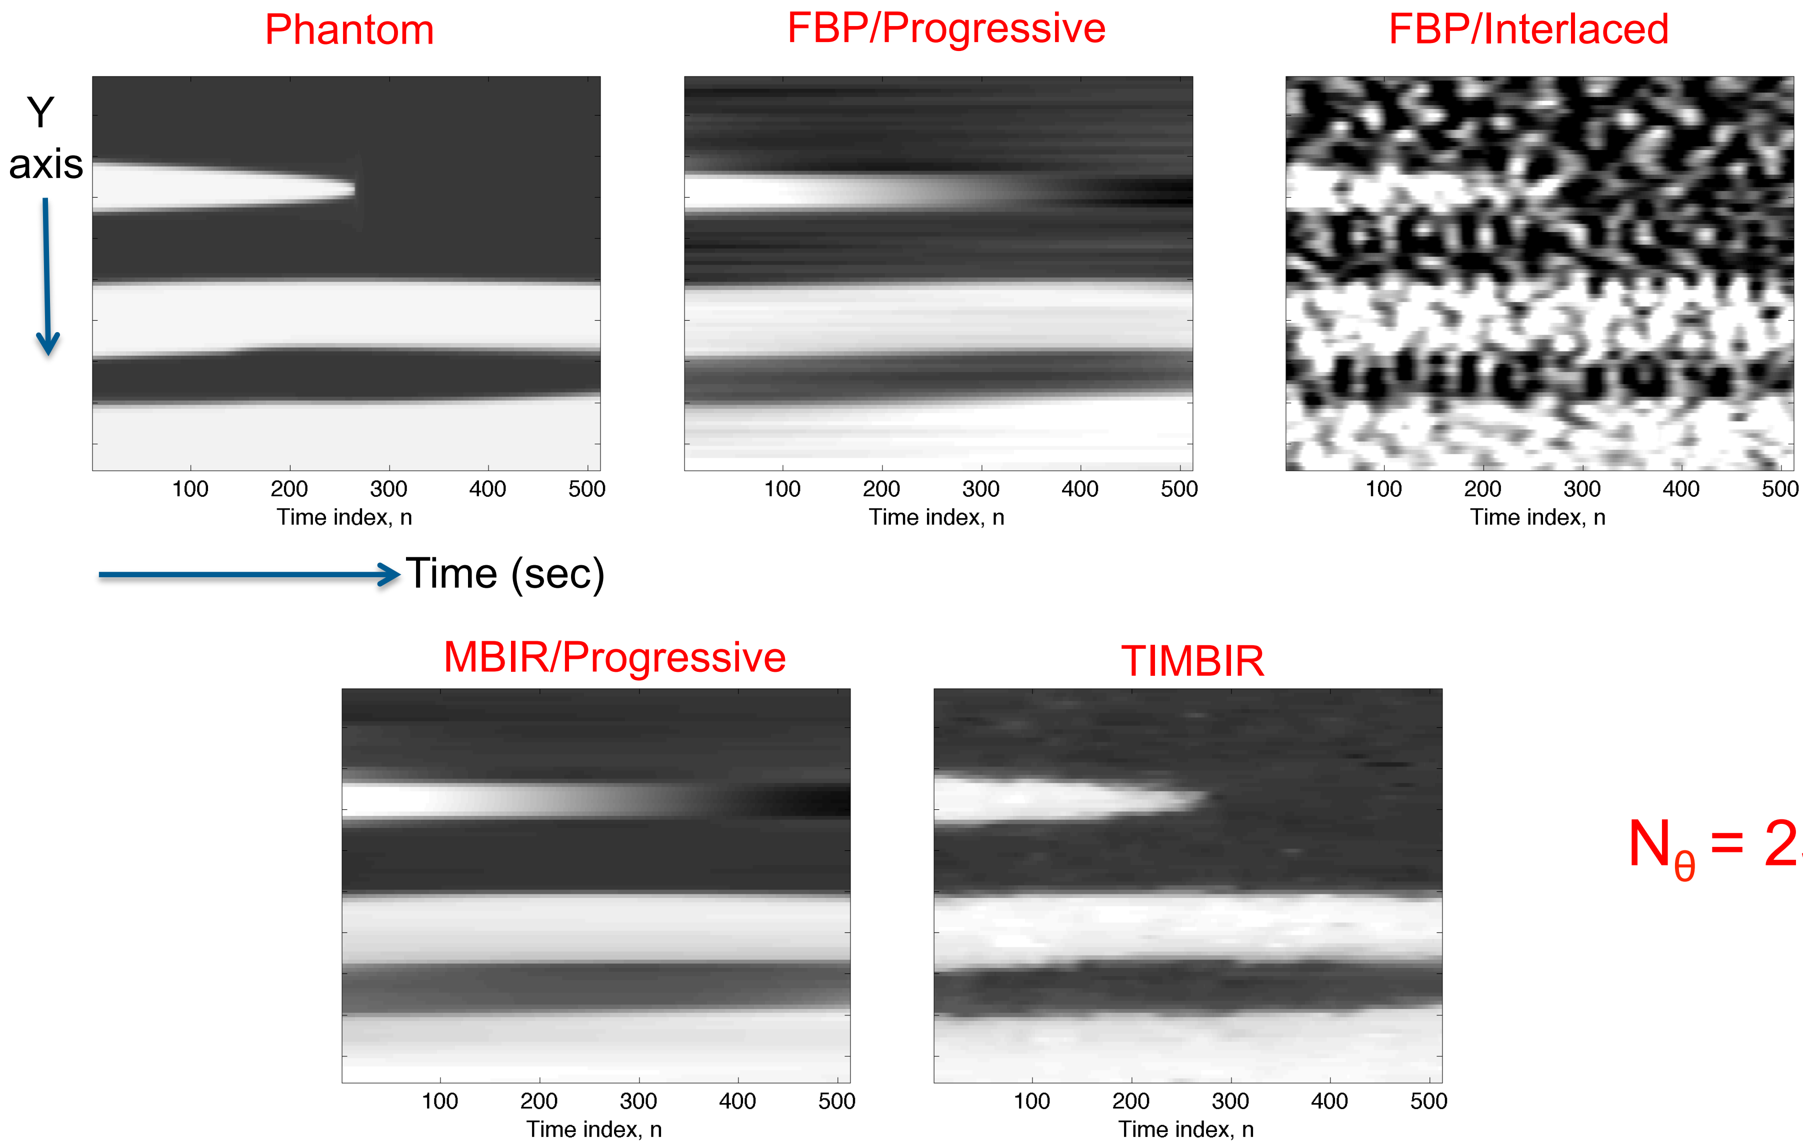

$N_\theta = 256$

TIMBIR: Synergy Between View Sampling and MBIR Reconstruction

- TIMBIR results in synergistic improvement

	FBP	MBIR
Conventional View Sampling	• Low noise robustness • Low temporal resolution • Medium quality	• High noise robustness • Low temporal resolution • High quality
Interlaced View Sampling	• Low noise robustness • High temporal resolution • Low quality	• High noise robustness • High temporal resolution • High quality **TIMBER**

Validation using Real Data

- Facilities
  - Advanced Photon Source (APS) synchrotron at Argonne National Laboratory
  - High performance computer at Advanced Light Source (ALS) at Lawrence Berkeley Laboratory
- Objectives
  - To reconstruct the solidification of Al-Cu microstructures at high temporal resolution.
  - Evaluate TIBIR method.

TIMBIR with K = 16

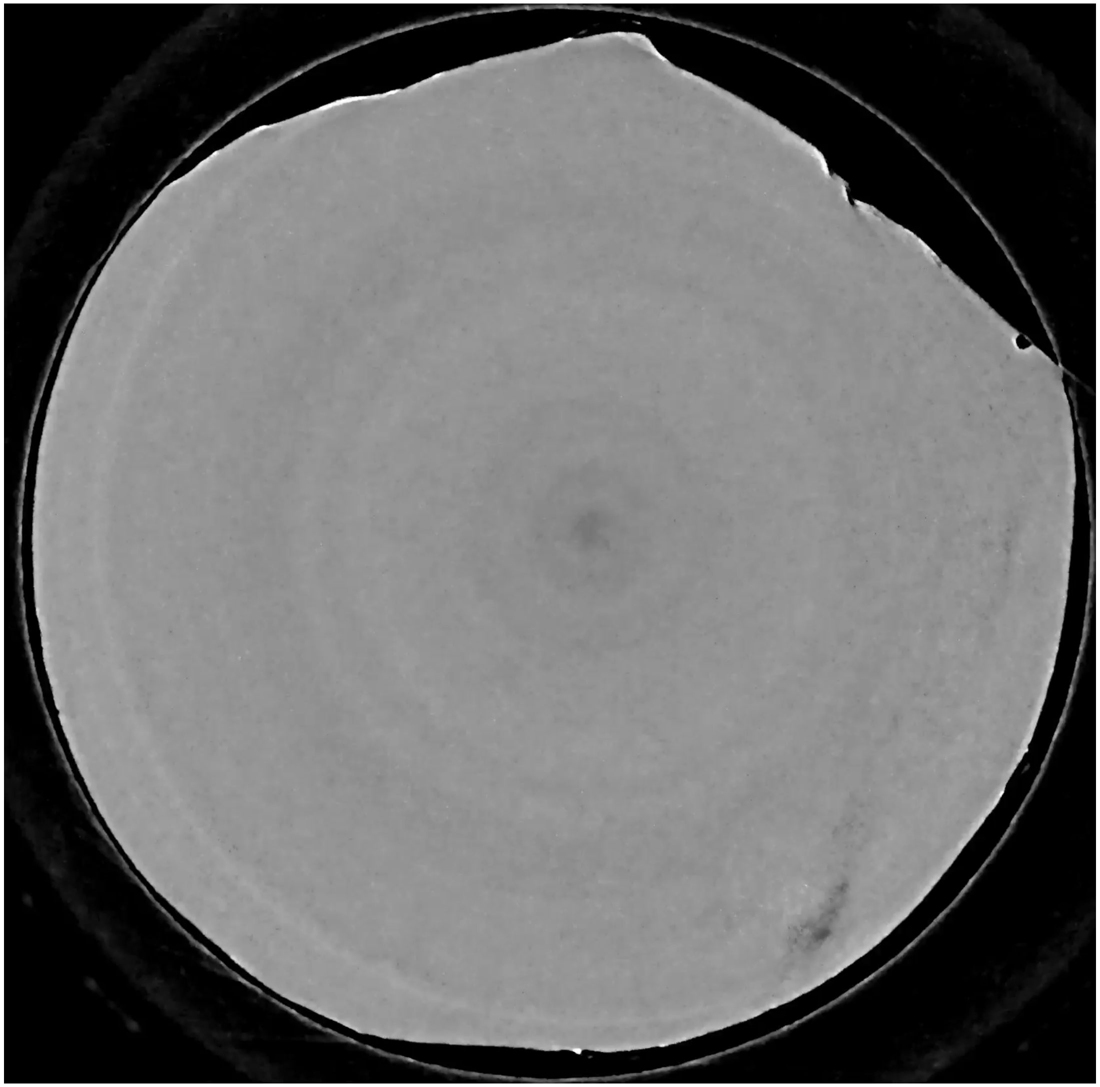

Single Spatial Slice

- Experiment
  - Solidification of aluminum and copper mixture
  - Temperature decreased at 2^0 Celsius per minute
  - k=16; r=16; N_v=2000
  - 16x speed up
- Reconstruction
  - (2048 x 2048 x 1000) space x 16 time
  - $(0.65\ \text{mm})^3$ voxel size
  - 1.8 sec time step
  - Image scaling:10000 HU to 60000 HU

4D Segmentation of TIMBIR with K = 16

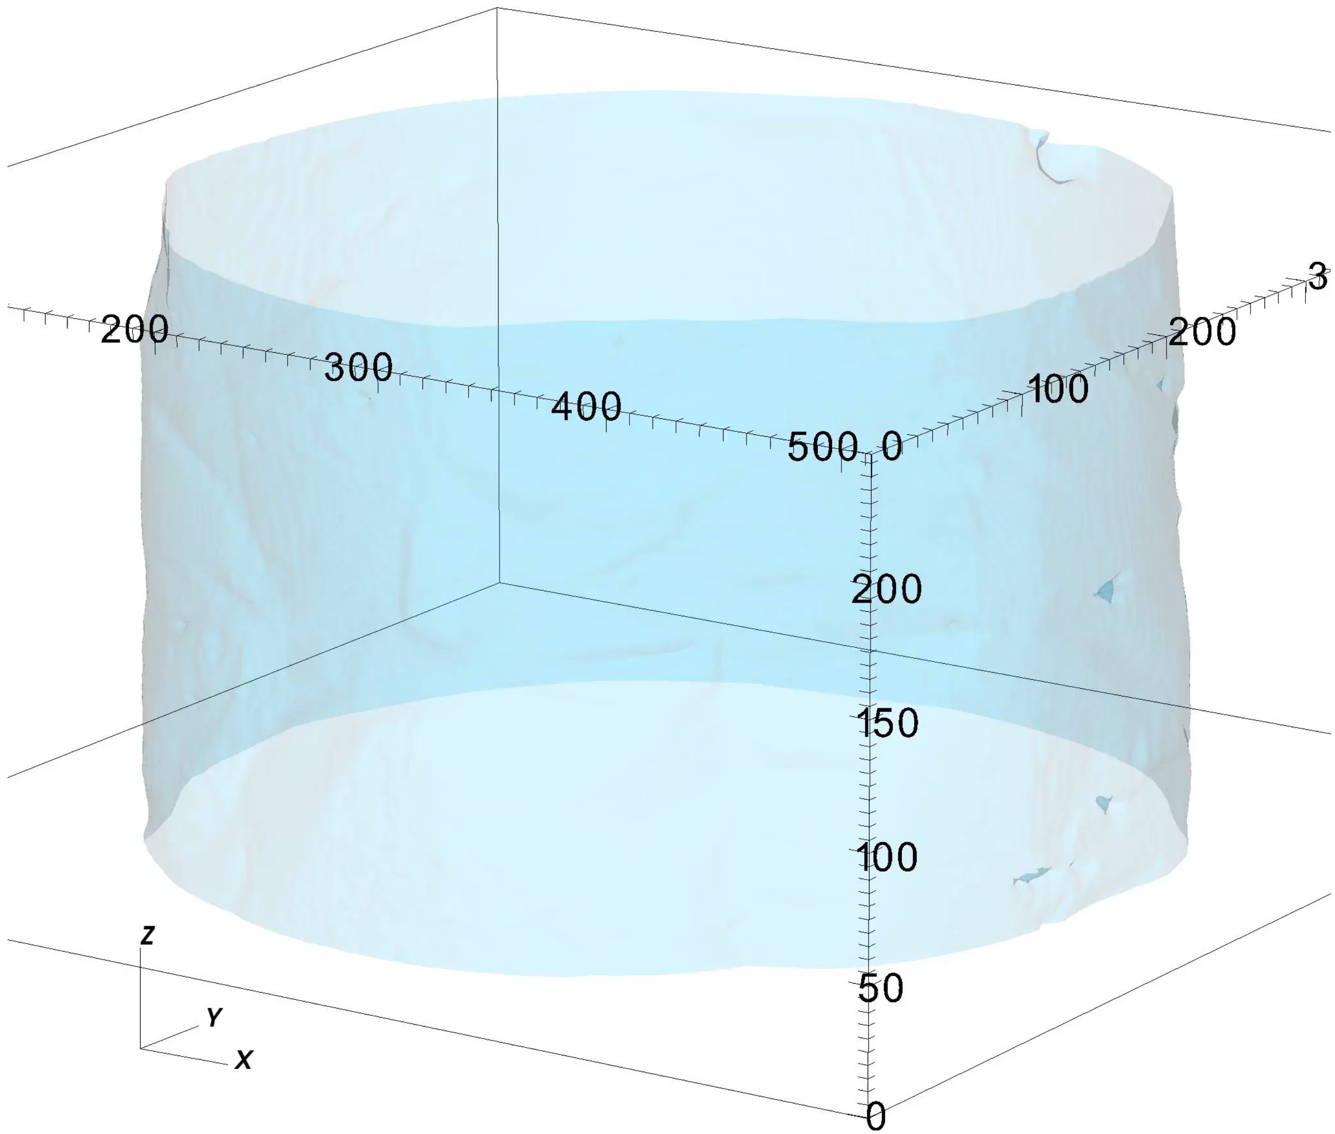

Electron Microscopy (EM) Microscopy for Material Science

Venkat Venkatakrishnan, Purdue
Larry Drummy, AFRL
Marc De Graef, CMU
Jeff Simmons, AFRL

Electron Microscopy (EM) Imaging

- 2-D Characterization of samples (biology, material science)
- Various modalities (Bright Field, Dark Field etc.)

STEM

Bright Field

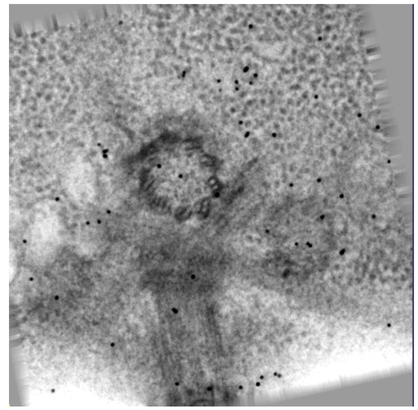

Biological sample*

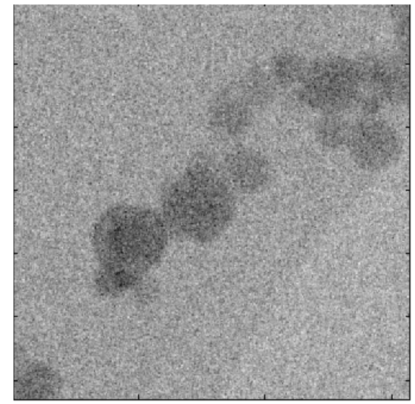

Aluminum nanoparticles**

Dark Field

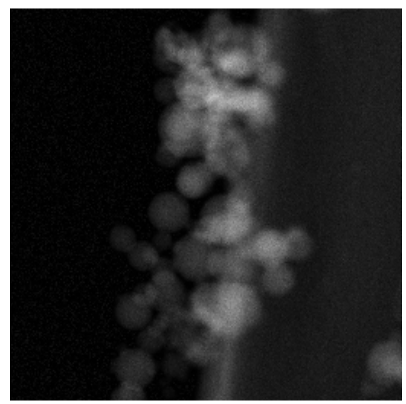

Aluminum nanoparticles**

*http://bio3d.colorado.edu/imod/doc/etomoTutorial.html
** L.F. Drummy, AFRL

Bright Field (BF) vs. Dark Field Imaging

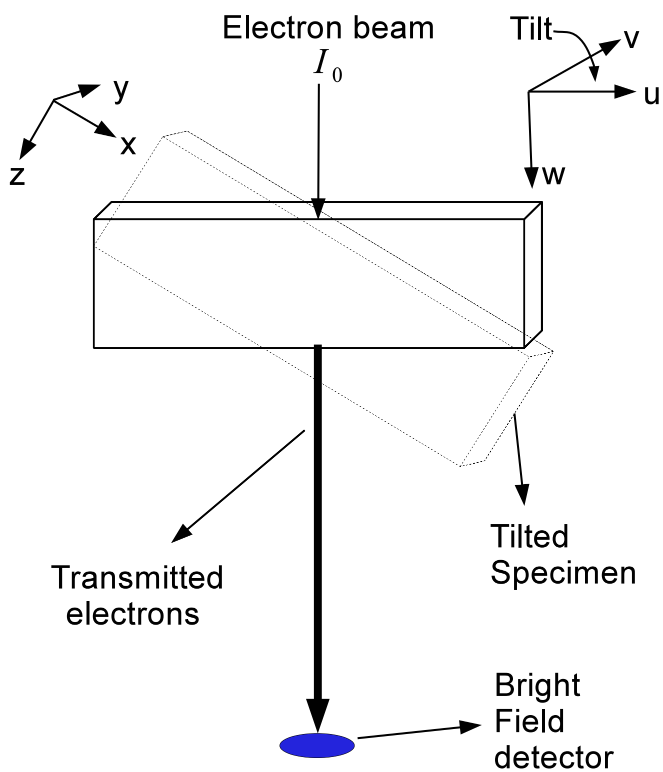

Bright Field:

Image is **bright** when sample is removed

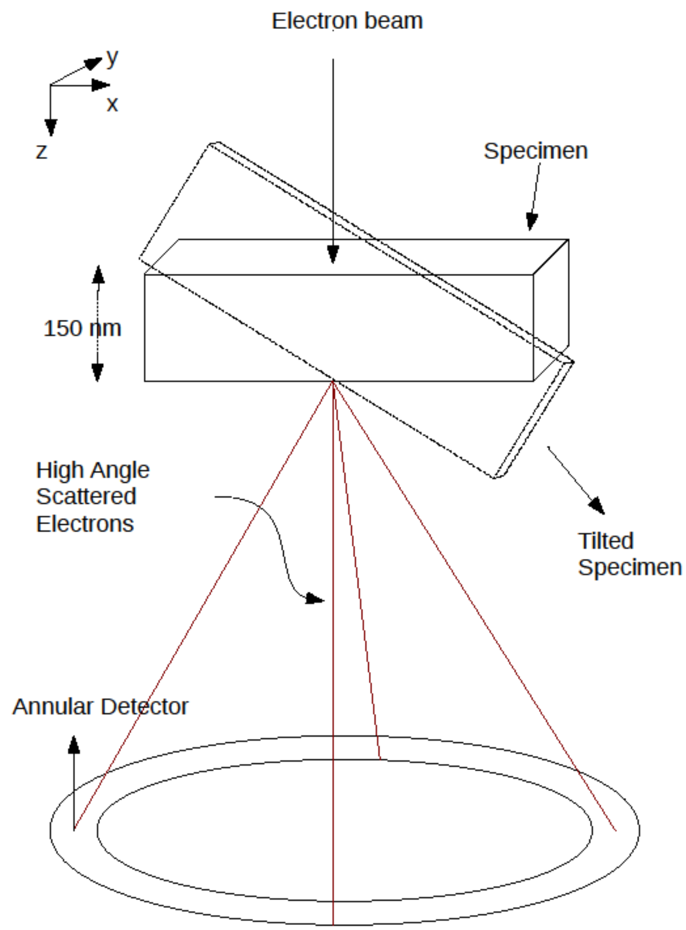

Dark Field:

Image is **dark** when sample is removed

The Problem with Bright Field EM

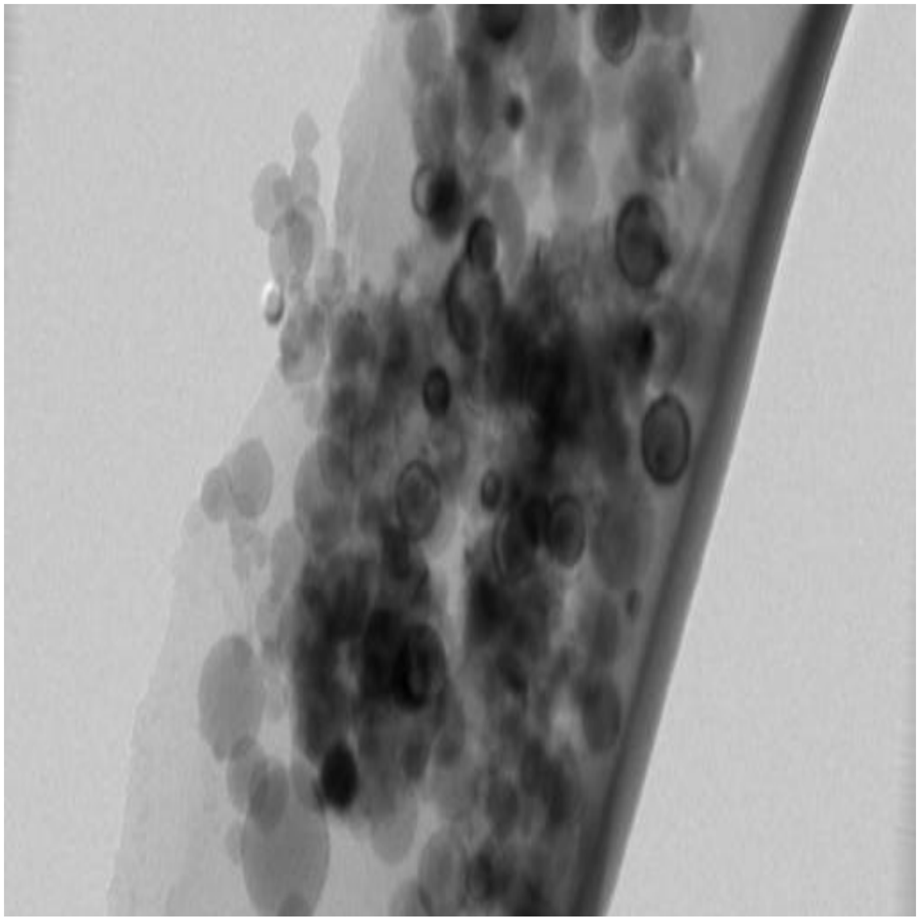

Aluminum nano particles

- Crystalline materials create Bragg scatter
- When Bragg scatter occurs, particle is dark => Beers Law is wrong!

$$\int_{ray} \mu(r)\, dr \neq -\log\left(\frac{\lambda_j}{\lambda_0}\right)$$

- **“Tomography doesn’t work”**

MBIR Reconstruction with Bragg Rejection

$$\left(\hat{f}, d\right) = \arg\min_{f \geq 0, d} \left\{ \underbrace{\frac{1}{2} \sum_{i=1}^{M} \beta_{T,\delta}\left(\left(g_i - A_{i*} f - d\right)\sqrt{\Lambda_{ii}}\right)}_{\text{Forward Model With Bragg Rejection}} + \underbrace{\sum_{\{i,j\} \in \chi} w_{ij} \rho\left(f_i - f_j\right)}_{\text{Prior Model}} \right\}$$

f : Linear attenuation coefficients to reconstruct (nm^{-1})

$g_i = -\log(\lambda_i)$

$d = -\log(\lambda_D)$

λ_i : Measured BF signal (counts)

λ_D: Unknown Dosage (counts) - can be estimated

$$\beta_{T,\delta}(e) = \begin{cases} e^2 & |x| \leq T \\ 2T\delta|e| + T^2(1-2\delta) & |x| > T \end{cases}$$

→ **Eliminate the effect of Bragg anomalies**

Λ_{ii} : $\dfrac{1}{\text{Noise variance}}$ (scaled) for measurement i

A_{i*} : i^{th} row of forward projection matrix

M : Total number of measurements

Generalized Huber Function

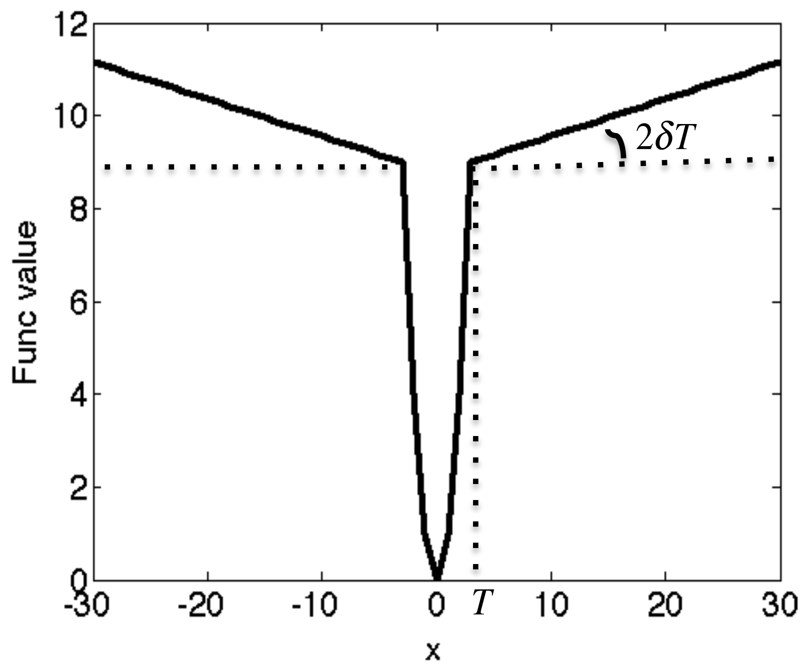

$$\beta_{T,\delta}(e) = \begin{cases} e^2 & |x| \le T \\ 2T\delta|e| + T^2(1-2\delta) & |x| > T \end{cases}$$

$e \rightarrow$ Measurement error

$\beta_{T,\delta} \rightarrow$ Generalized Huber function

- Reduce effect of outliners due to Bragg
- Generalized Huber Function
  - Proper distribution => ML estimation of threshold T
  - Surrogate function (majorization) for optimization

Al –TEM Data (Movie)

Reconstruction params

Region : $Y - [920,974]$

$X \ - [563,1484]$

Thickness = 350 nm

$p = 1.2$

$\sigma_f = 1.6 \times 10^{-4} \text{nm}^{-1}$

Recon Voxels = $(2 \times 0.83)^3$ nm^3

$T = 3; \delta = 0.5$

Region used for reconstruction

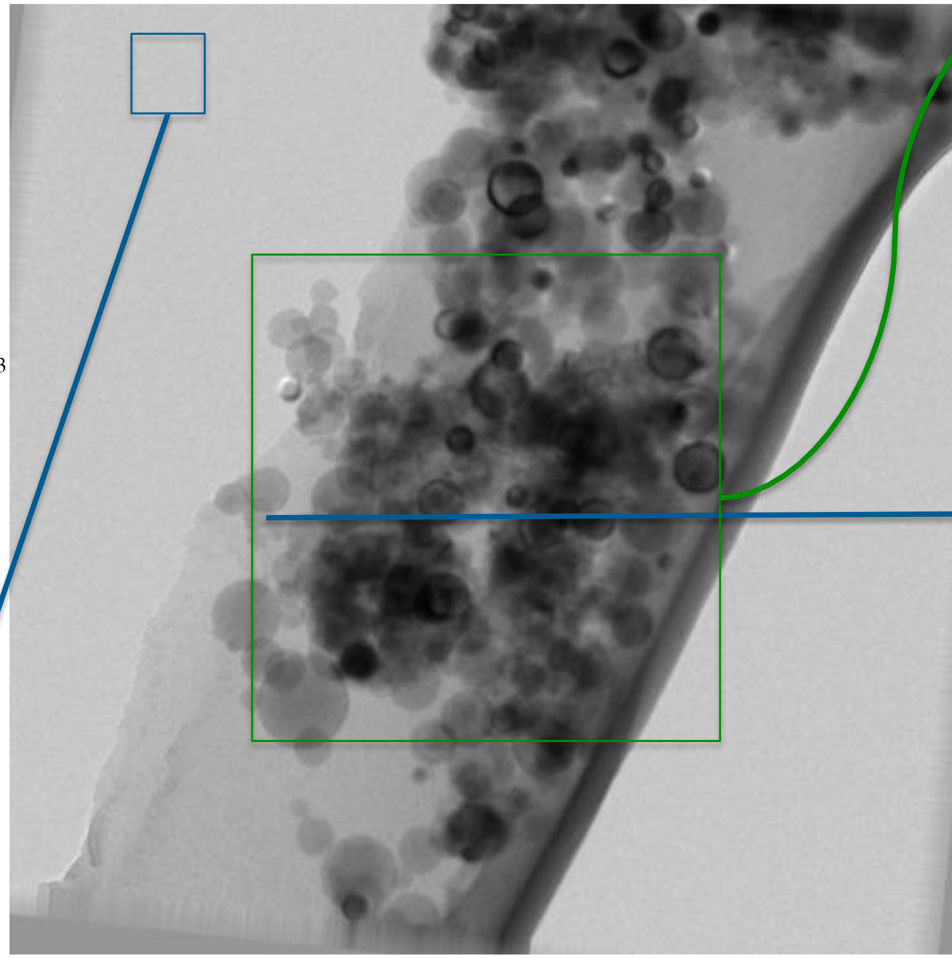

Approximate location of displayed slice

Average value in a void region

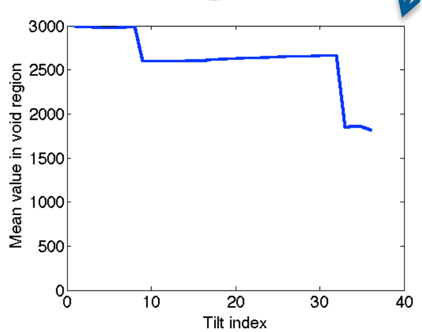

Data range (int) : [-32728,-21780]

Preprocess to (uint) : [40,10988]

Reconstruction : x-z cross section

FBP*

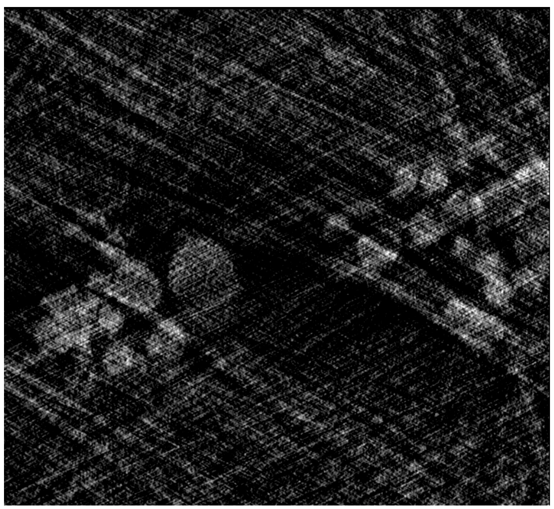

MBIR – No anomaly correction

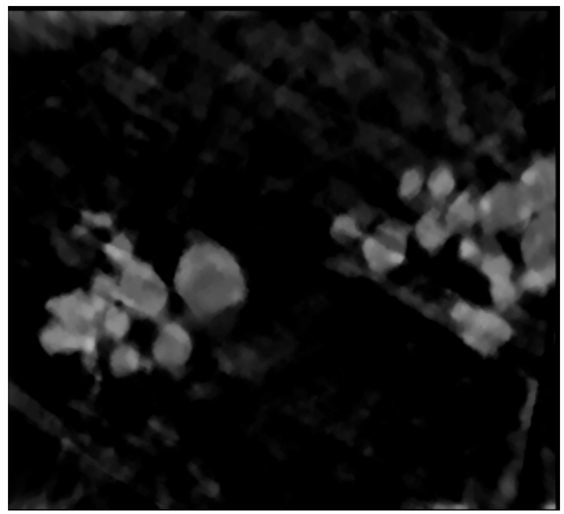

MBIR – with anomaly correction

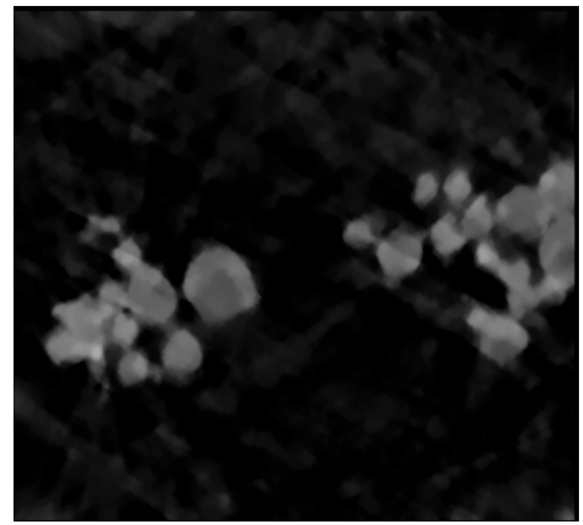

$\sigma_f = 1.6 \times 10^{-4} nm^{-1}$

Reconstruction : x-y cross section

FBP*

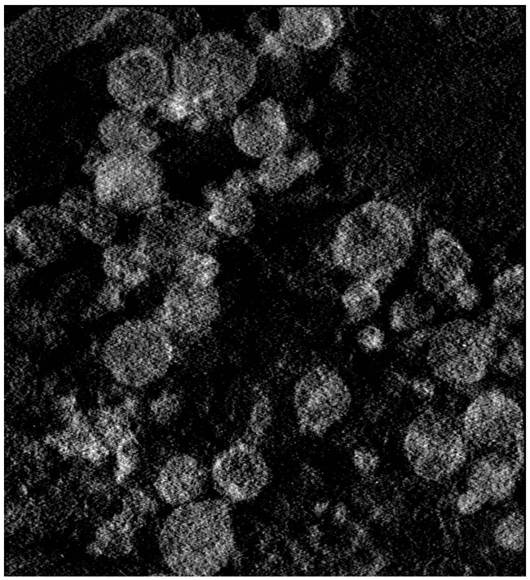

MBIR – No anomaly correction

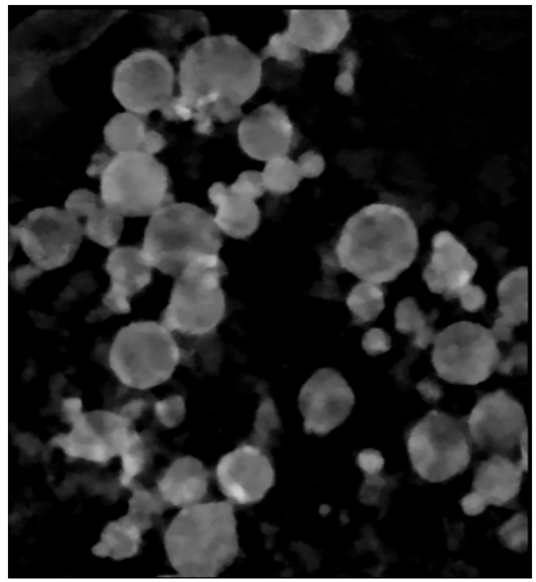

MBIR – with anomaly correction

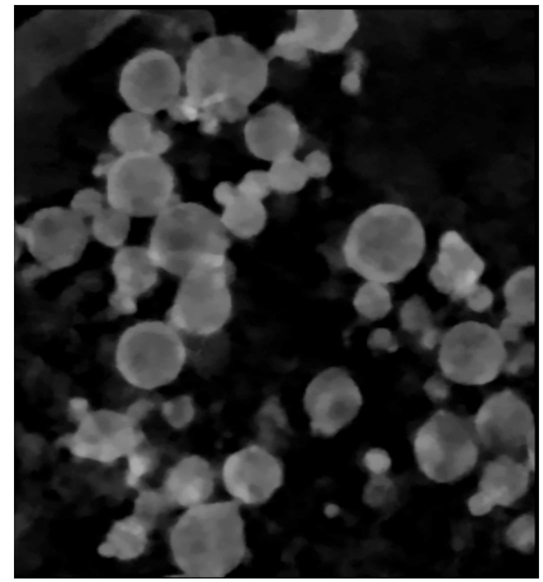

Bragg Anomaly Classification

sinogram

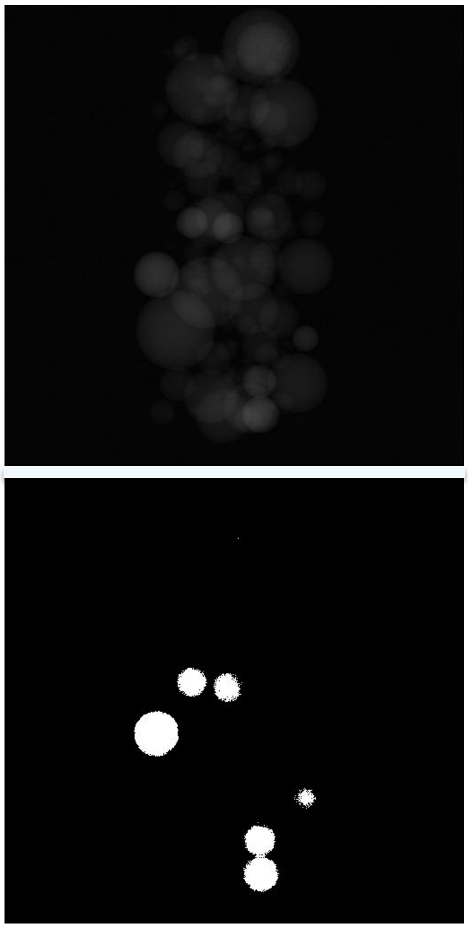

anomaly classifier

- MBIR-BF cost function labels Bragg
- Identify Bragg signature of particles
- Requires 3D segmentation

$\hat{\sigma}^2 = 3.03$

Fraction classified : 3.92%

“Bragg Feature Vector” Extraction Algorithm

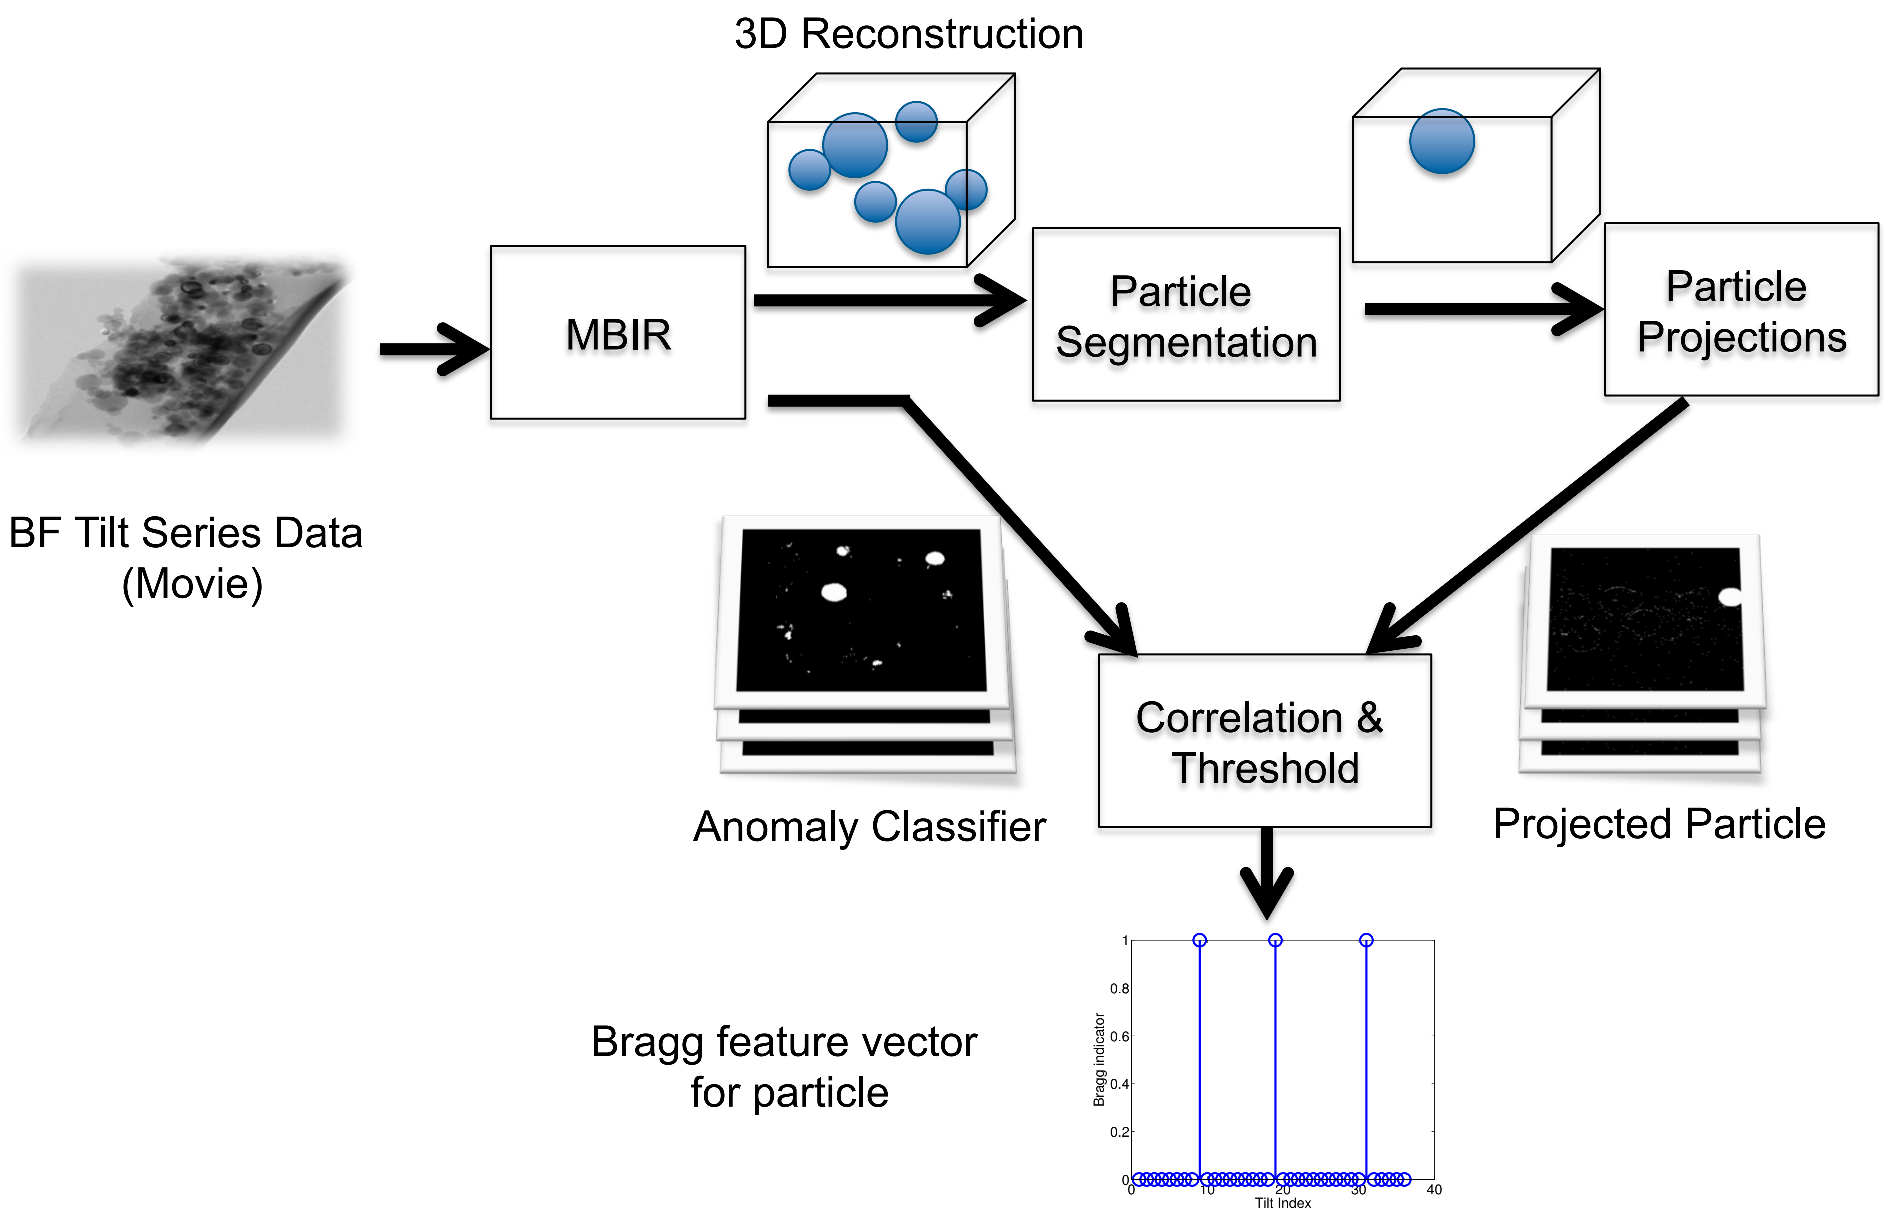

Extracted Bragg Feature Vectors

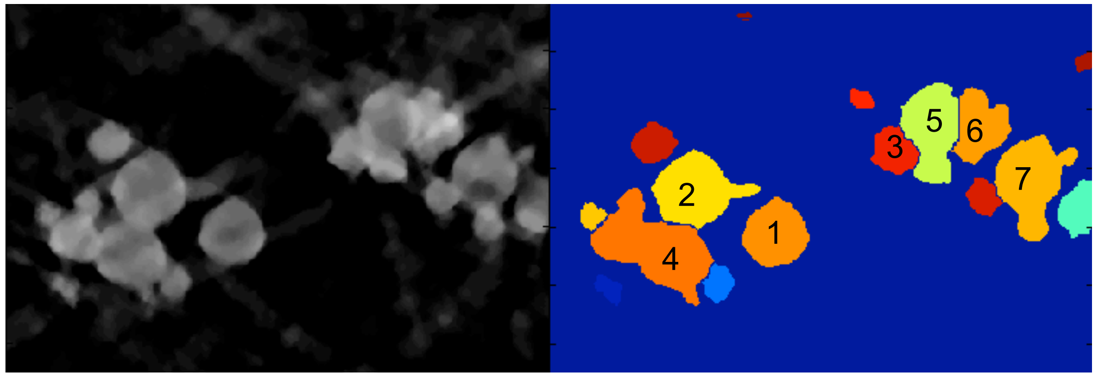

MBIR Reconstruction Segmented Particles

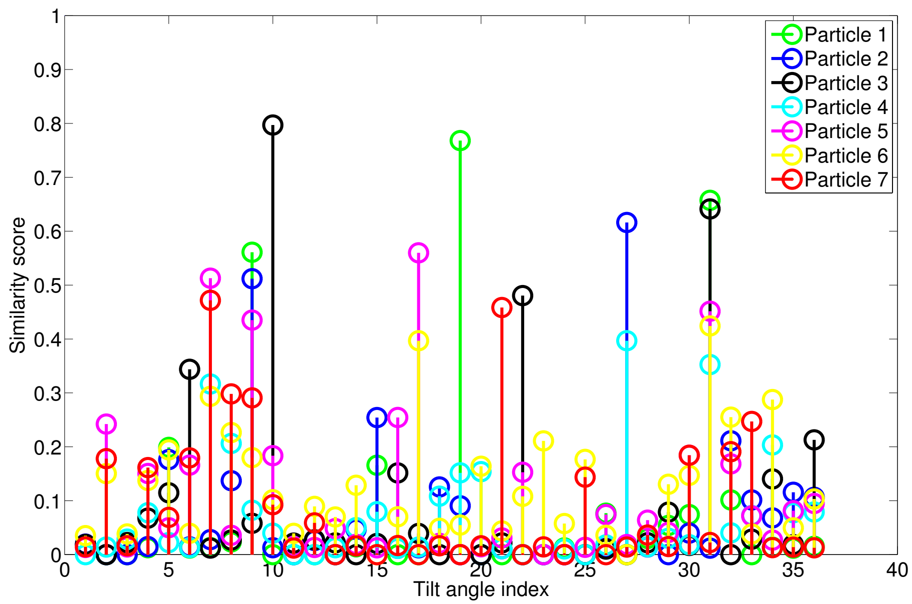

Bragg Feature Vector for Each Particle

Advanced Priors Models

S. Venkat Venkatakrishnan, Purdue University
Brendt Wohlberg, Los Alamos National Laboratory
Suhas Sreehari, Purdue University
Garth J Simpson, Purdue University
Charles A. Bouman, Purdue University

Prior Modeling of Images

- An open problem of great importance
  - Low, mid, high level models
  - Crucial in denoising problems
- Promising recent approaches:
  - MRFs; Dictionary bases learning methods; kSVD; Non-local means; BM3D; Bilateral filters; Gaussian mixture models (GMMs)
- Many of these are not really prior models

Plug & Play Priors Algorithm

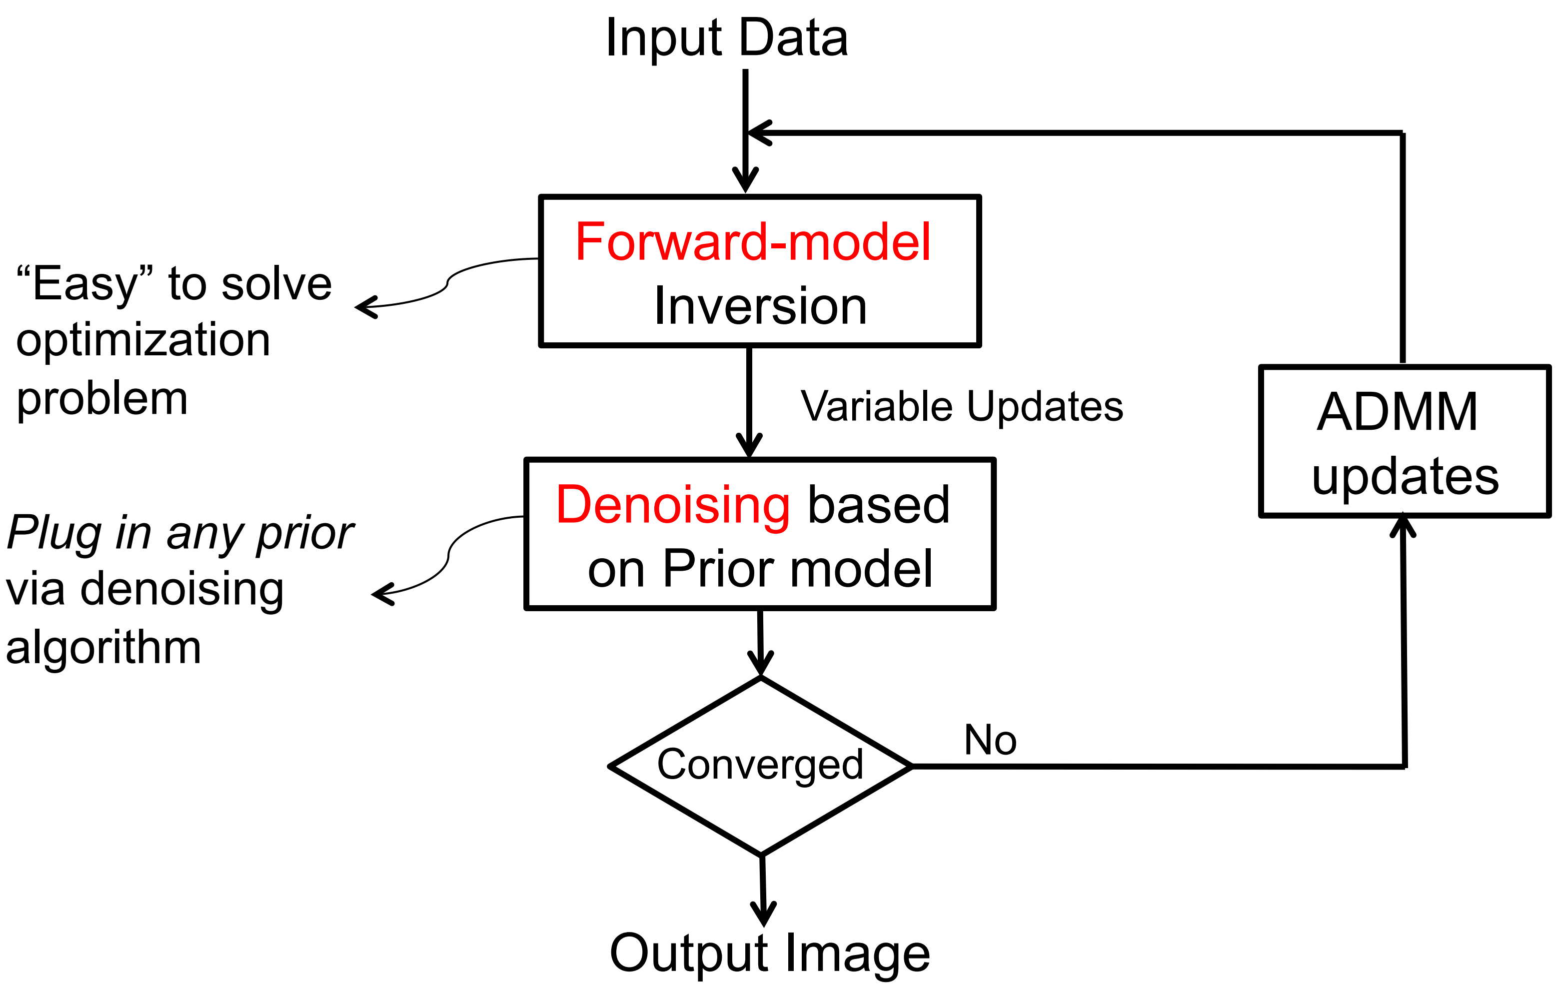

- **Can use "any" denoising model as a "prior"**

Deblurring with Many “Priors”

Ground Truth

Blurred, Noisy Data

K-SVD

RMSE : 13.13

BM3D

RMSE : 13.91

PLOW

RMSE : 14.95

TV

RMSE : 15.21

q-GGMRF

RMSE : 14.14

Inpainting with Many “Priors”

Ground Truth

Subsampled Image

Noise std. dev : 5% of max signal

K-SVD

RMSE : 14.11

BM3D

RMSE : 12.56

PLOW

RMSE : 14.54

TV

RMSE : 15.50

q-GGMRF

RMSE : 15.72

Space/Time Scanning Optical Microscope

Garth Simpson, Purdue University

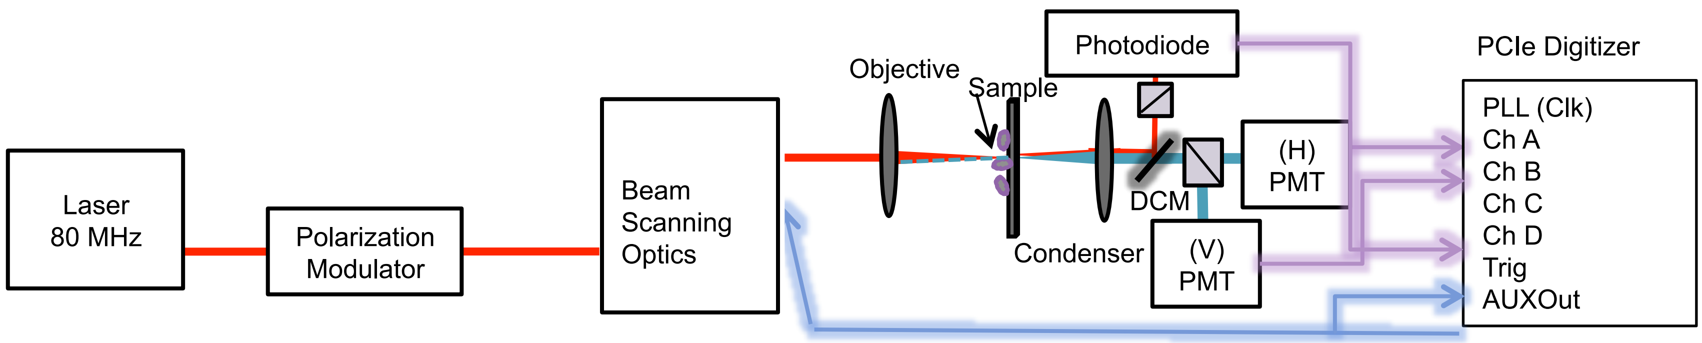

- Scan dynamic sample in space and time

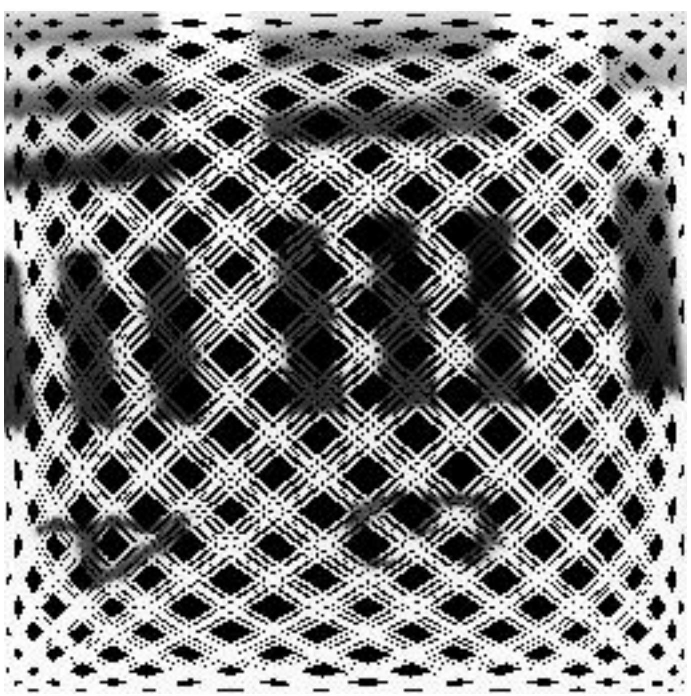

Original time slices

Inpainted time slices

STEM Inpainting on Real Data*

Larry Drummy, AFRL

Ground Truth*

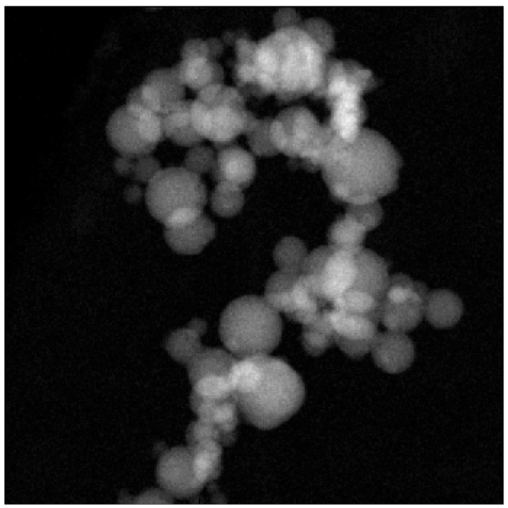

20 % of pixels uniformly sampled

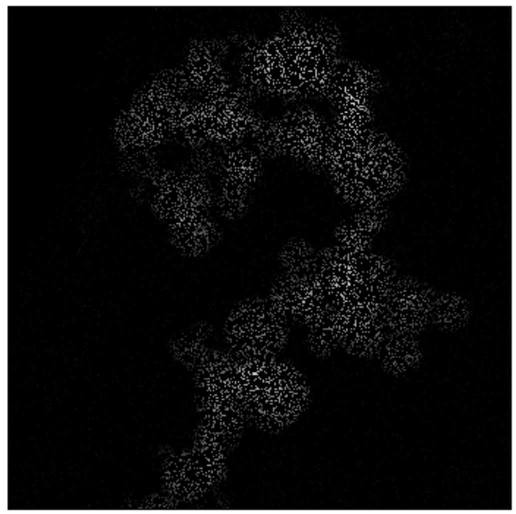

Plug-and-Play reconstruction with **BM3D** “prior”

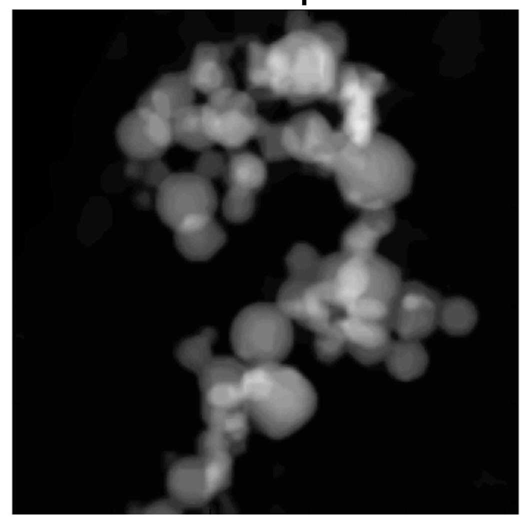

* Cropped out a region from real data and rescaled

Model Based Dynamic Sampling (MBDS)

- *Dilshan P. Godaliyadda, Purdue University*
- *Prof. Gregery T. Buzzard, Purdue University*
- *Prof. Charles A. Bouman, Purdue University*

Raster Sampling: Optimally Bad

- Each new sample provides **the least** information.

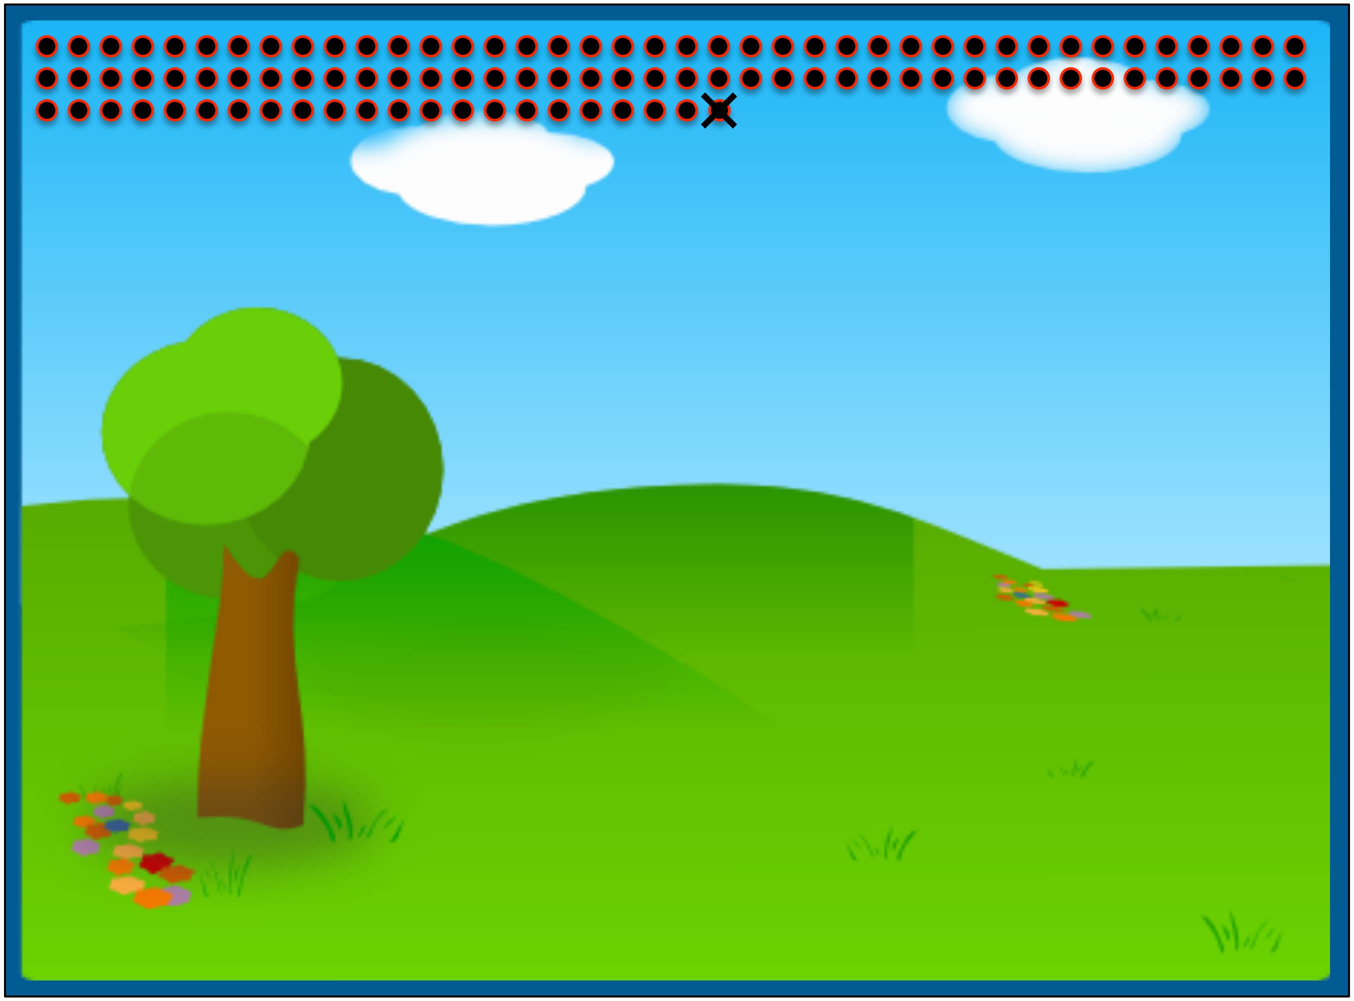

Random Sampling: Better

- Each new sample provides **much more** information.

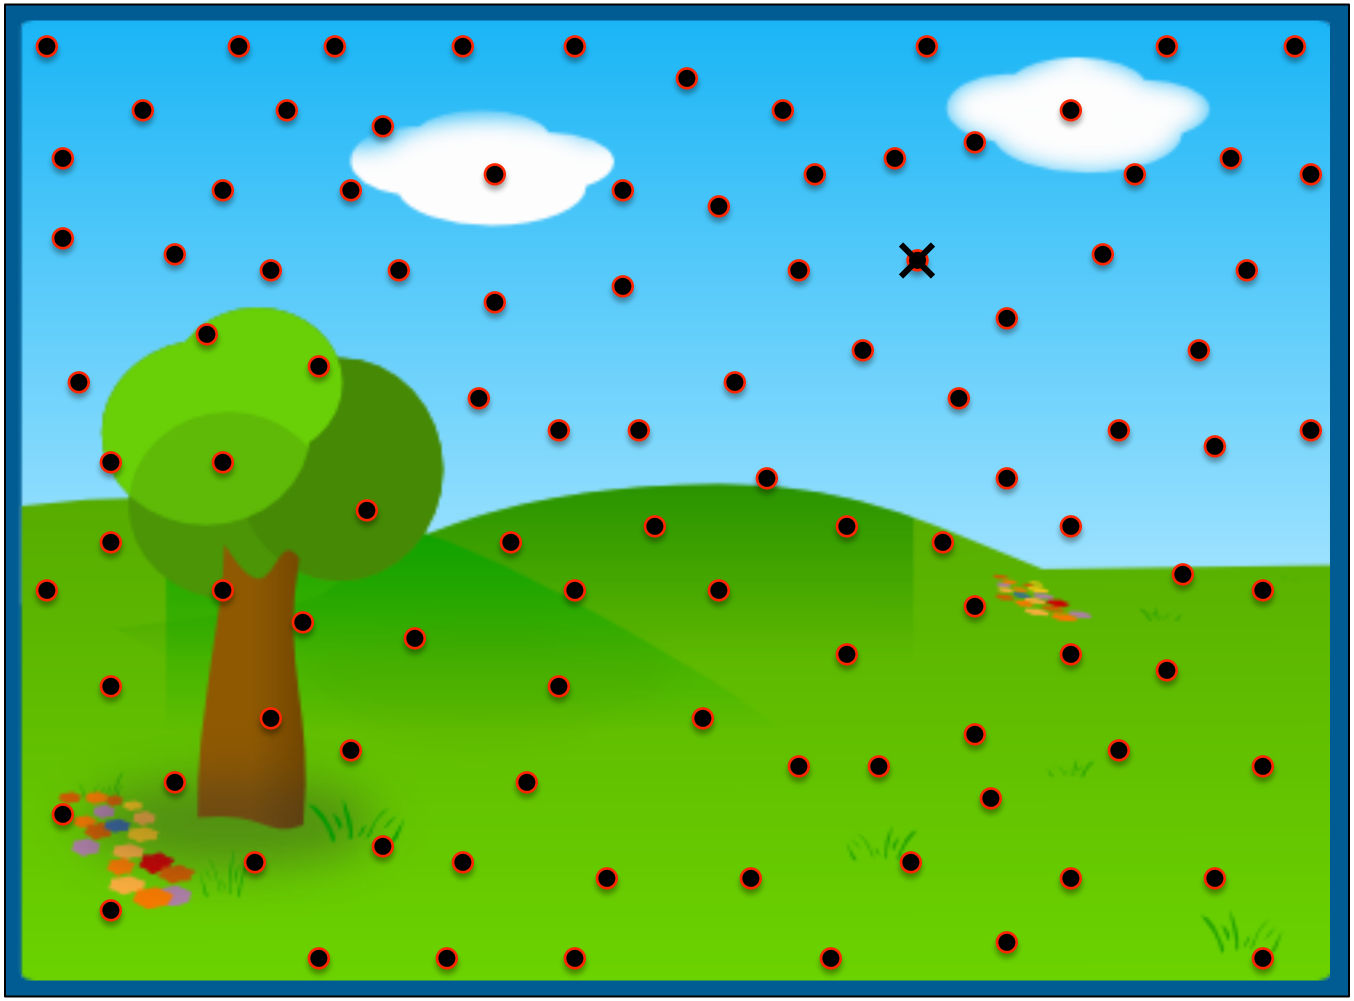

Optimal (Greedy) Sampling: Best

- Each new sample provides **the most** information.

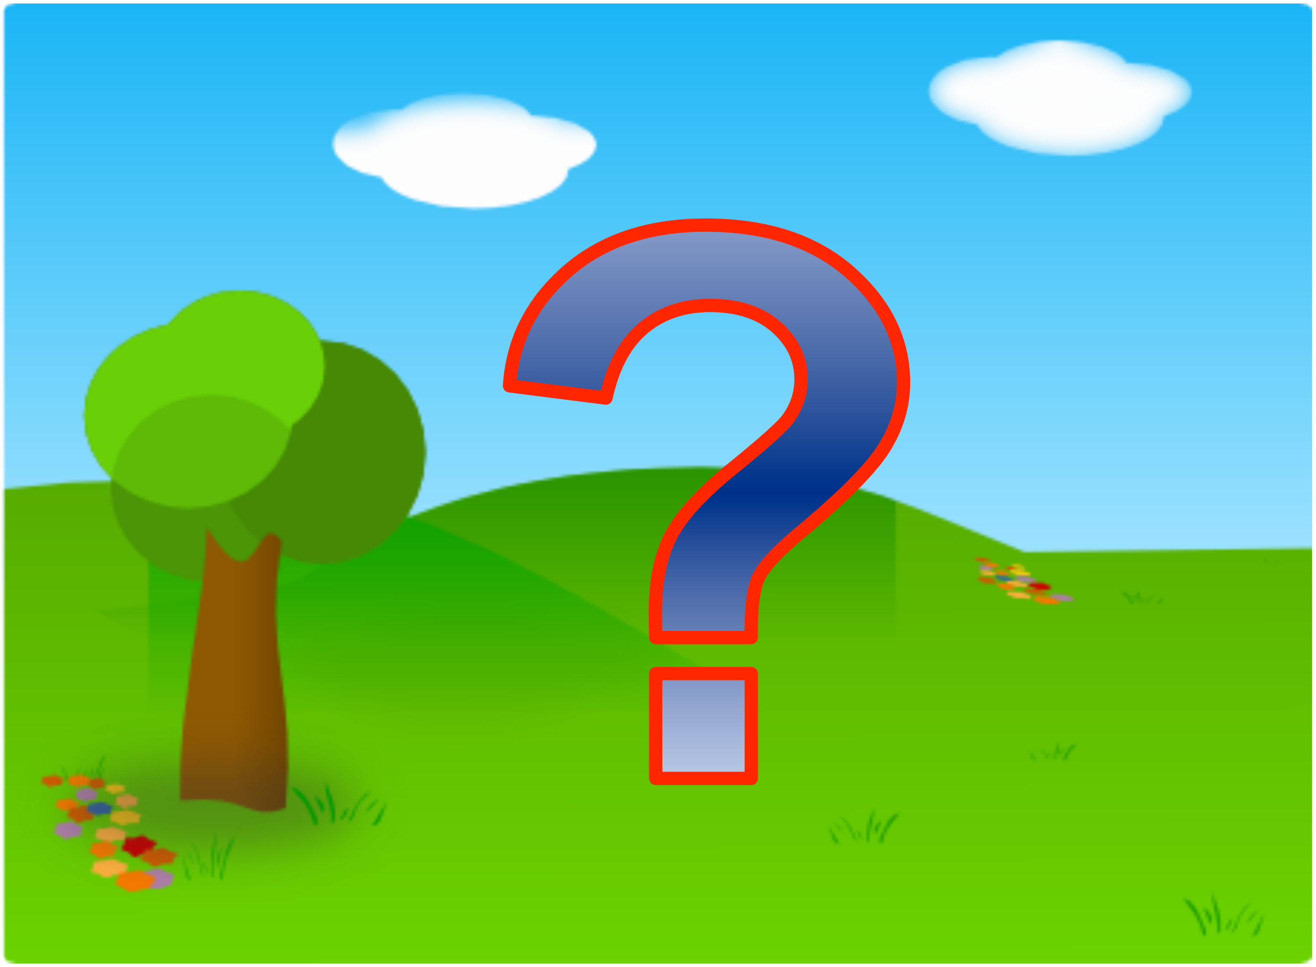

Recursion for Optimal Greedy Sampling

For each new sample {

$$y^{(k)} = A^{(k)}x + w^{(k)}$$

Step 1: Measure signal

$$\mu_{x|y}^{(k)} = \mathrm{E}\left[x \middle| y^{(k)} \right]$$

$$R_{x|y}^{(k)} = \mathrm{E}\left[\left(x - \mu_{x|y}^{(k)}\right)\left(x - \mu_{x|y}^{(k)}\right)^T \middle| y^{(k)} \right]$$

Step 2: Find posterior covariance of image

When everything is Gaussian,

$R_{x|y}^{(k)}$ does not depend on $y^{(k)}$!!!!

$$m^{(k)} \leftarrow \arg\max_{m \in M} \left(m^T R_{x|y}^{(k)} m \right)$$

Step 3: Select pixel with largest variance

$$A^{(k+1)} = \begin{pmatrix} A^{(k)} \\ m^{(k)} \end{pmatrix}$$

Step 4: Add new row to A matrix

}

Non-Gaussian Prior => Dynamic Sampling

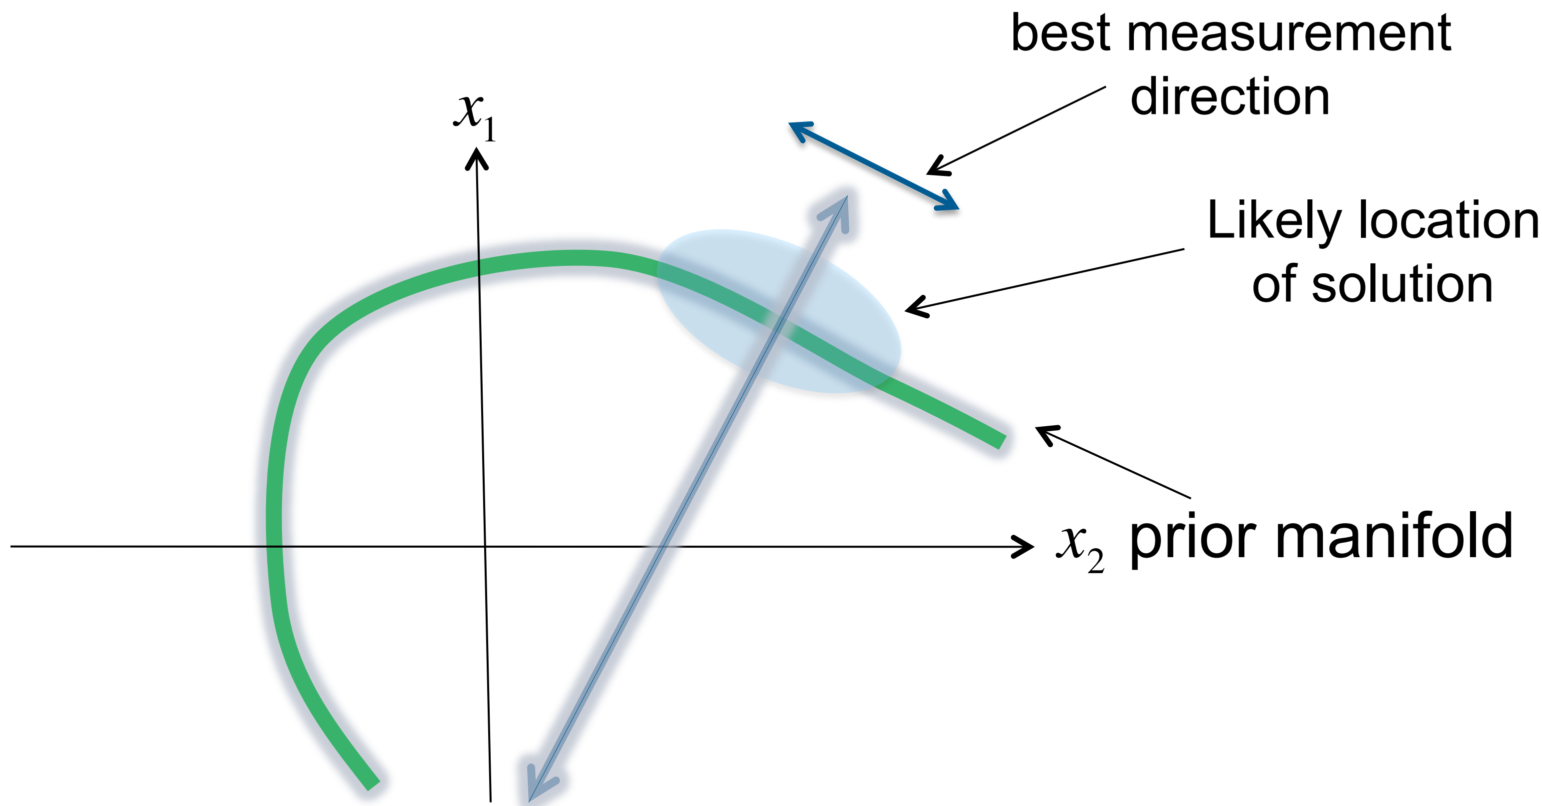

Non-Gaussian Prior => Dynamic Sampling

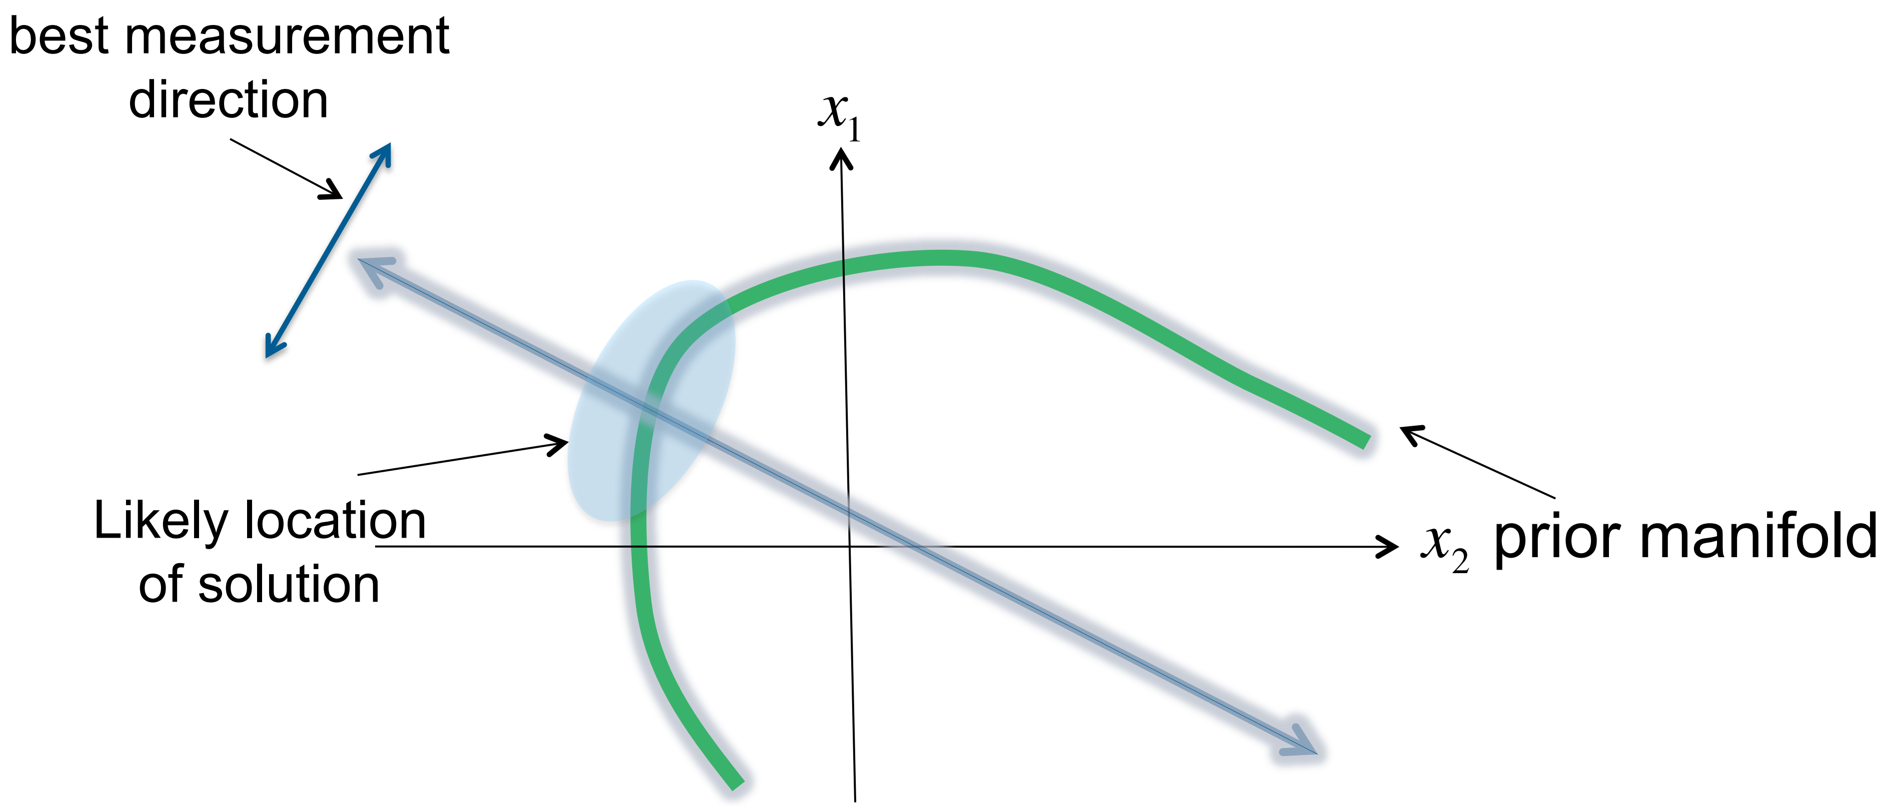

- Optimal sampling depends dynamically on previous samples

- **Non-Gaussian => Intractable calculation of posterior ☹**

Solution: Hastings-Metropolis Sampling of Posterior

For each new sample {

$$\boxed{y^{(k)} = A^{(k)} x + w^{(k)}}$$

Step 1: **Measure signal**

$$\boxed{\left\{x^{(1)}, x^{(2)}, \cdots, x^{(L)}\right\} \sim p\left(x \mid y^{(k)}\right)}$$

Step 2: **Generate *L* samples from posterior**

$$\hat{\mu}_n \leftarrow \frac{1}{L} \sum_{i=1}^{n} x_n^{(i)}$$

$$\hat{\sigma}_n^2 \leftarrow \frac{1}{L-1} \sum_{i=1}^{L} \left(x_n^{(i)} - \hat{\mu}_n\right)\left(x_n^{(i)} - \hat{\mu}_n\right)^T$$

Step 3: **Estimate posterior variance**

$$\boxed{n^{(k)} = \arg\max_n \left(\hat{\sigma}_n^2\right)}$$

Step 4: **Select pixel with largest variance**

$$A^{(k+1)} = \begin{pmatrix} A^{(k)} \\ 0, \ldots, \underset{\overset{\wedge}{n^{(k)}}}{1}, \ldots, 0 \end{pmatrix}$$

Step 5: **Add new row to A matrix**

}

Dynamic Sampling of 1D Signal

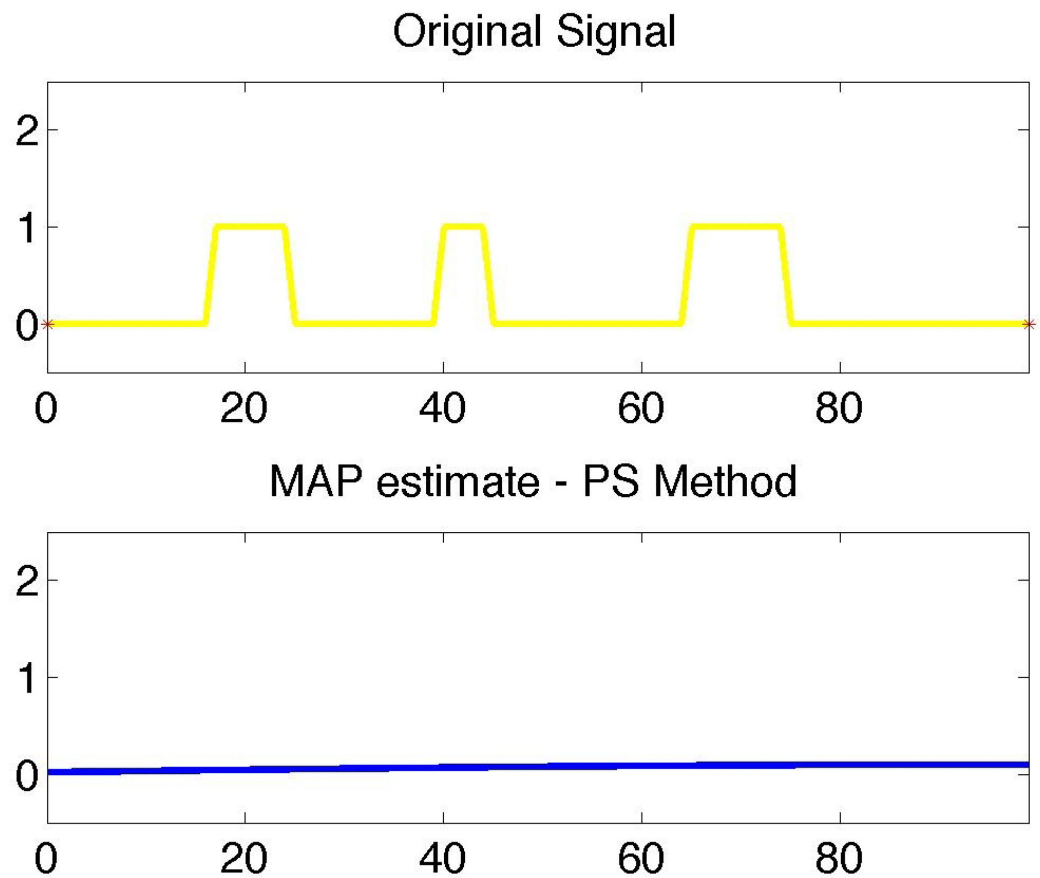

How to do this computation tractably in 2D?

2D Hastings-Metropolis Sampler

- Select pixel with largest sample variance

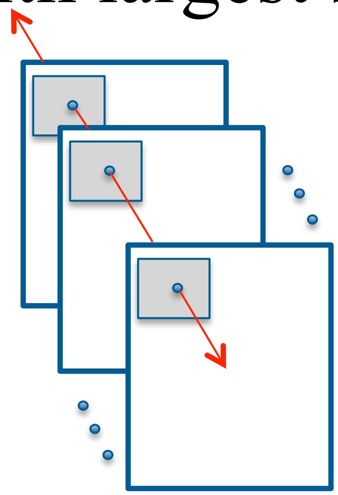

L=20 images generated from posterior

- Replace window of pixels using Hastings-Metropolis

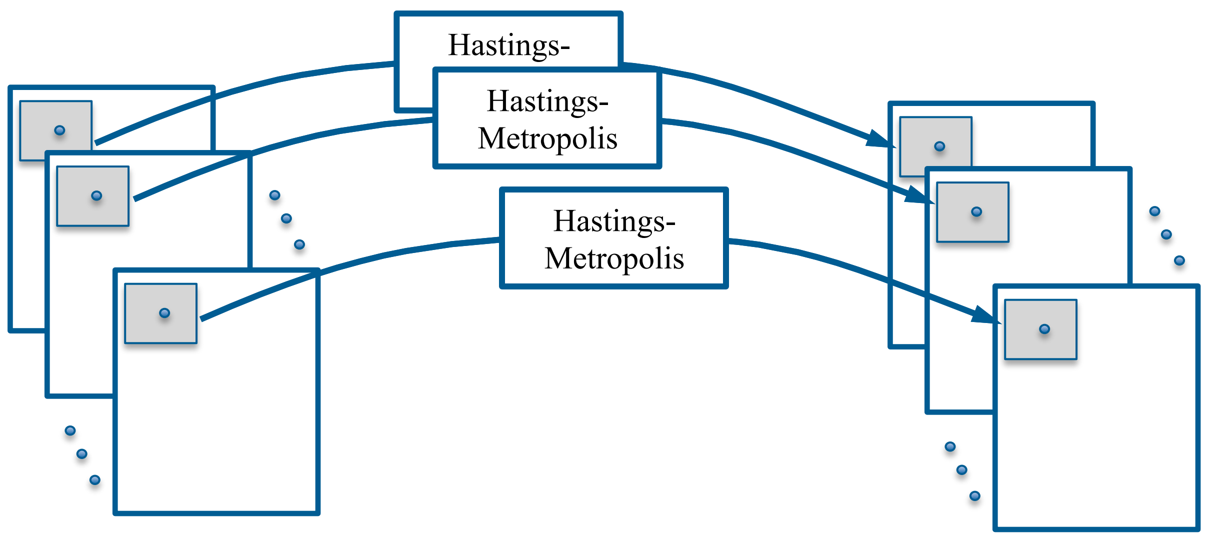

Dynamic Sampling for Phantom

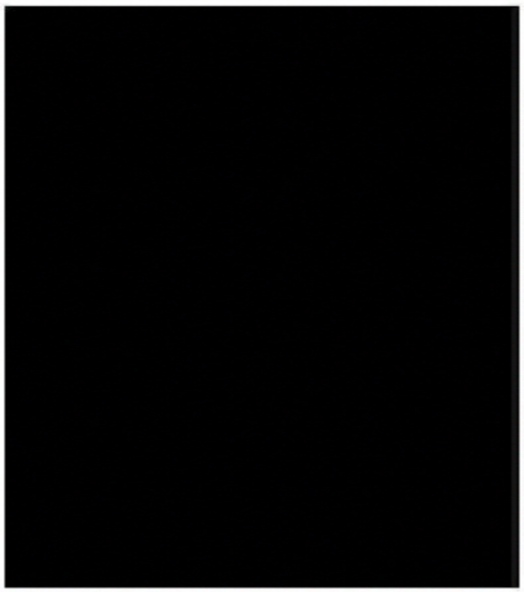

Selected Samples

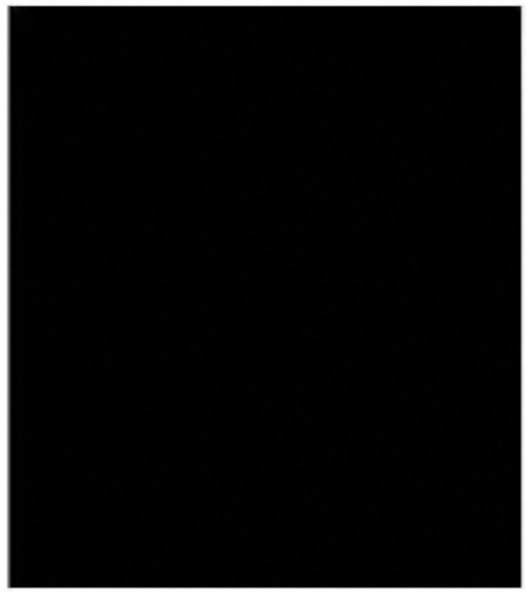

Measured Image

Reconstructed Image

Dynamic Sampling vs Random Sampling - 13% of Image Measured

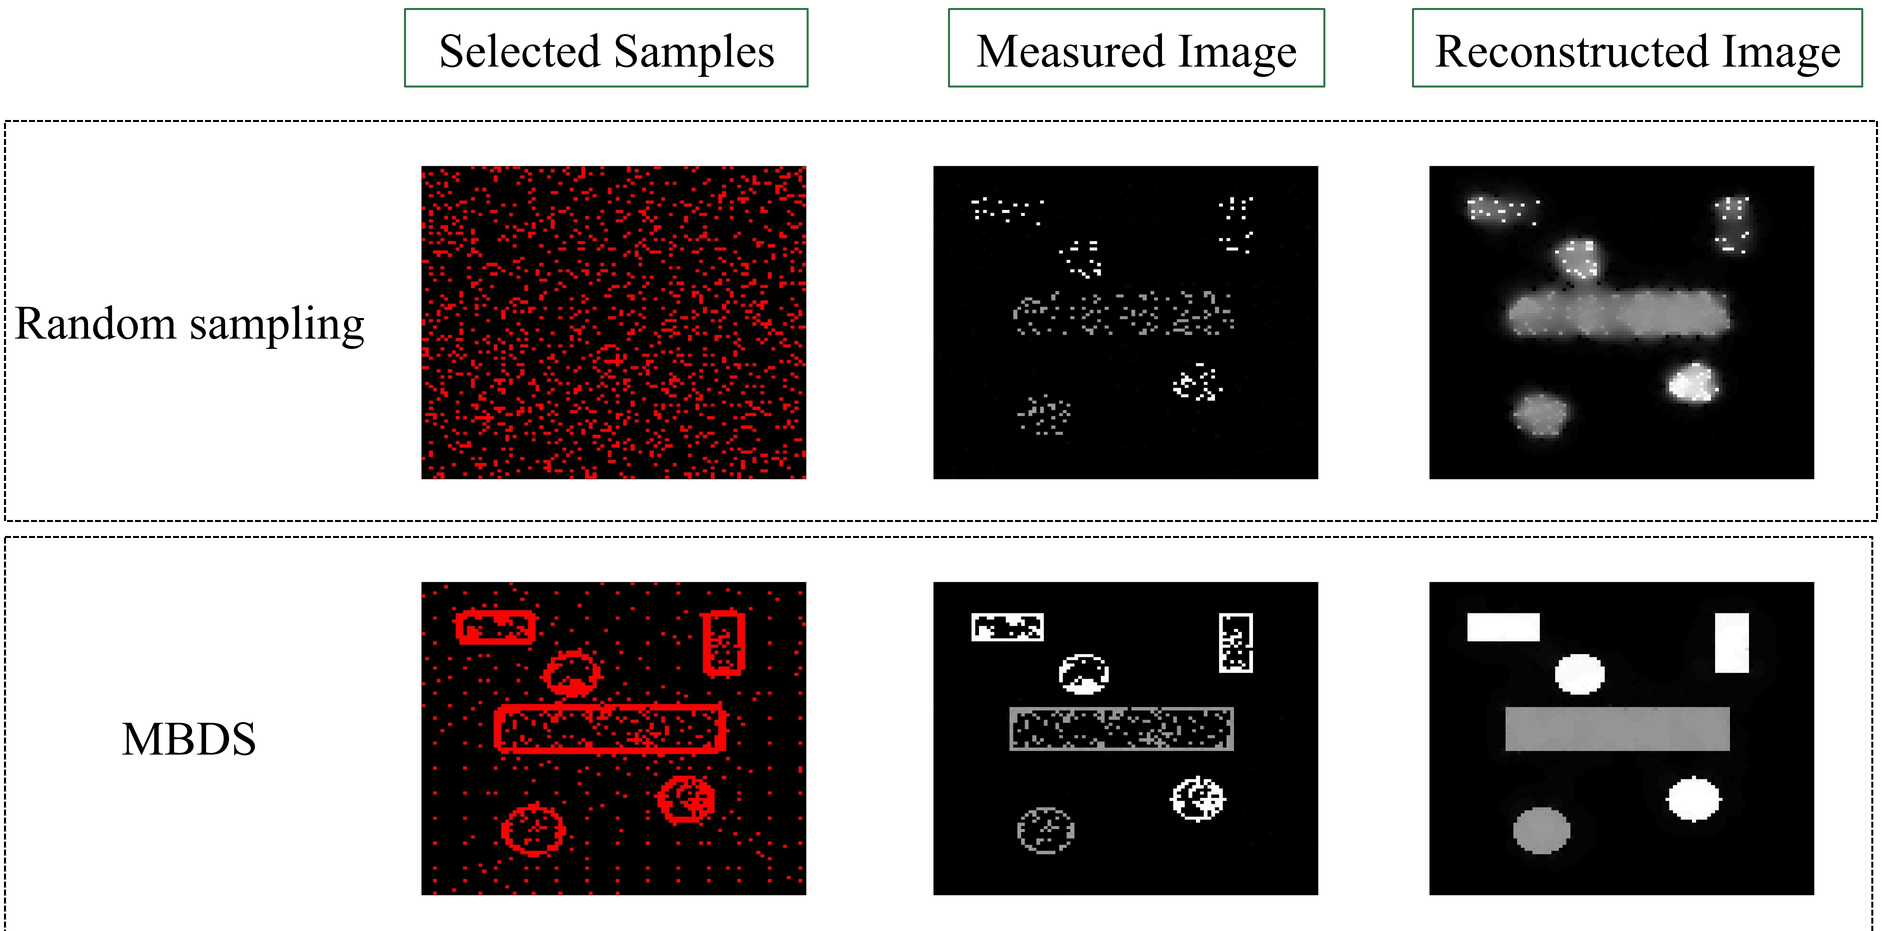

Dynamic Sampling for SEM Image

Selected Samples

Measured Image

Reconstructed Image

Dynamic Sampling vs Random Sampling - 8% of Image Measured

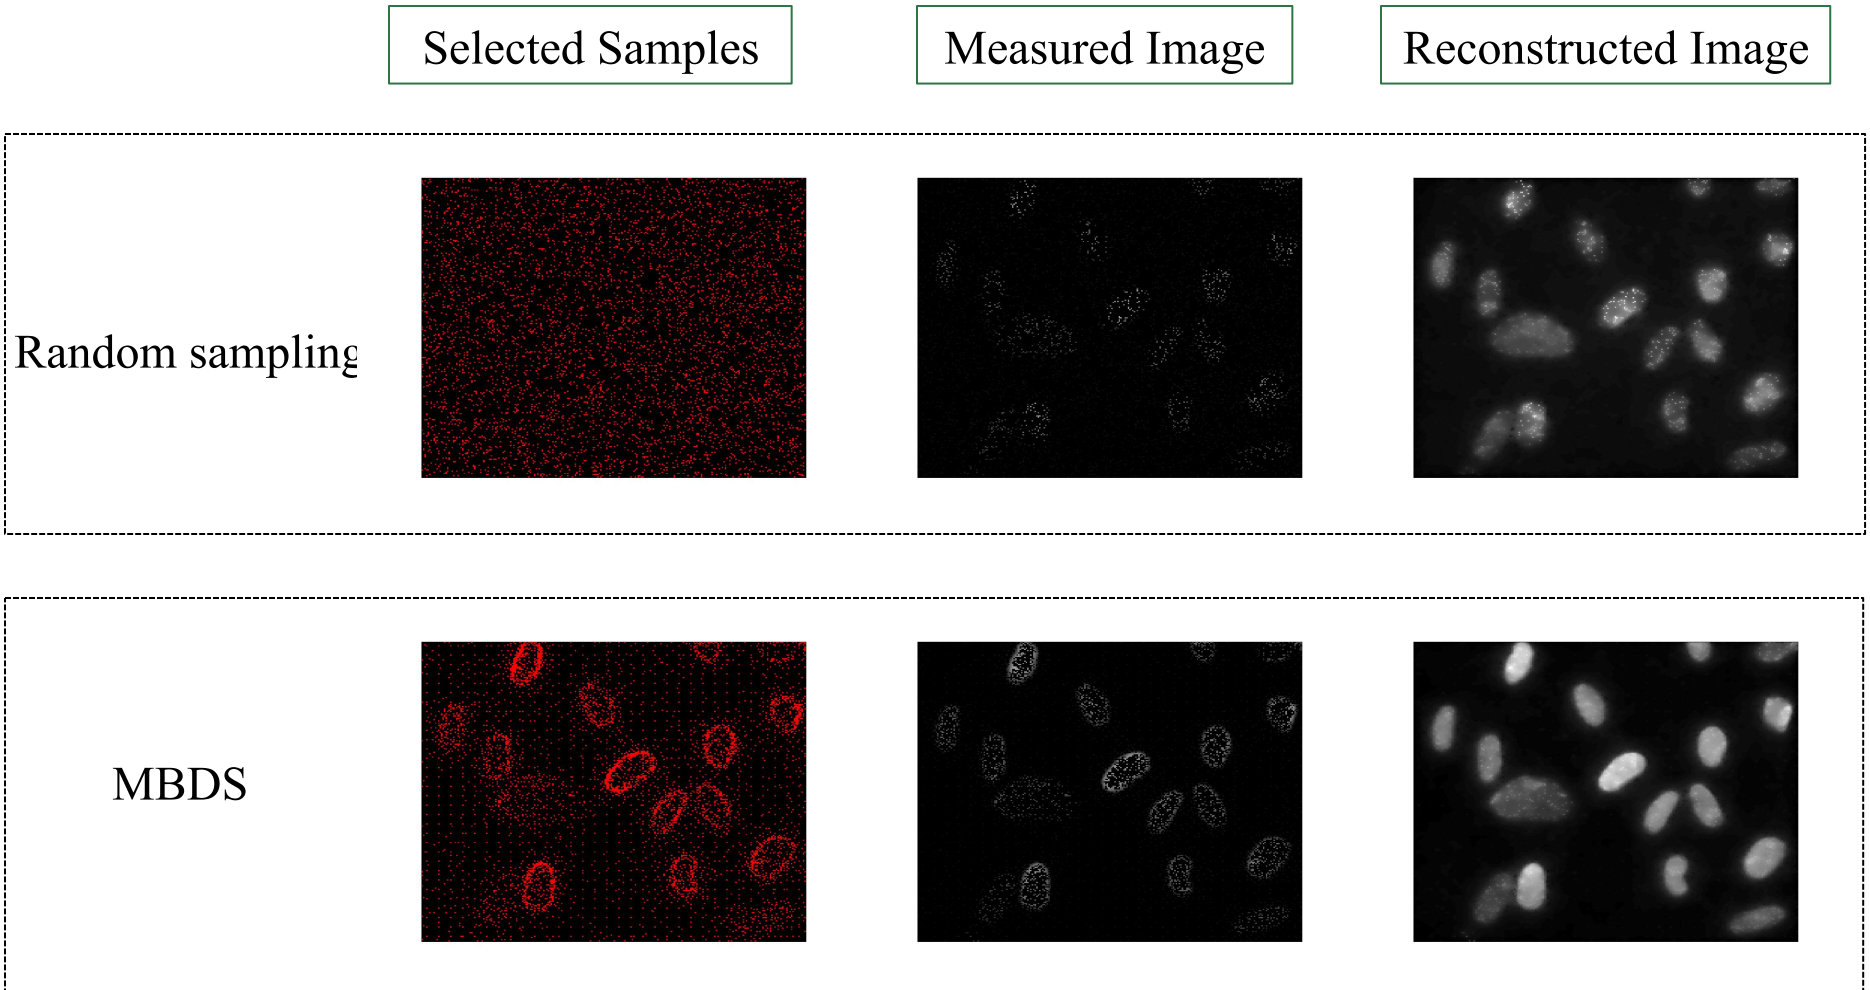

Analogy to Human Visual System

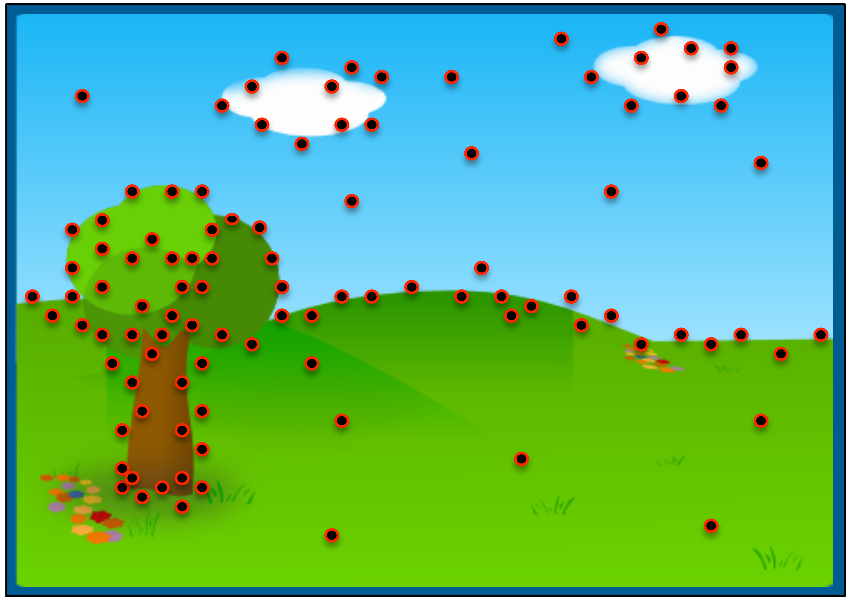

- Visual scanpath theory and saccadic movement of eye – *Stark, Privitera, Navalpakkam*
  - Bottom up => sensor model
  - Top down => prior model
- Interesting observations: Each pixel is selected to maximize information without knowledge of local edge structure.

Major directions for Integrated Imaging

- Creative design of sensor systems
- Image formation:
  - Forward modeling: Account for complex nonlinear parameters and models
  - Prior modeling: Account for properties of real images
- Community: Create interdisciplinary teams to solve high impact problems

- **New:** *IEEE Transactions on Computational Imaging*!