Bending Out of the Box: The Marriage of Sensors and Computational Imaging†

Charles A. Bouman, Purdue ECE/BME

2023 International Image Sensors Workshop

May 24, 2023

In collaboration with:

Greg Buzzard, Purdue Math
Soumendu Majee, Purdue ECE
Thilo Balke, Purdue ECE
Brendt Wohlberg, LANL
Craig Kemp, Eli Lilly
Singanallur V Venkatakrishnan, ORNL

†Thank you to Showalter Foundation, ORNL, LANL, NSF, GE Healthcare, AFRL, Eli Lilly, and DHS

Introduction to Computational Imaging

- What is Computational Imaging
- MBIR for Solving Inverse Problems
- Thin Manifold View of Inverse Problems

Imaging Over the Years*

silver halide

Daguerreotype
1839

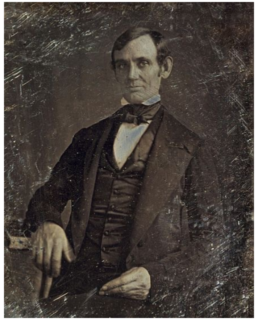

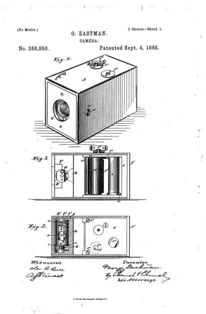

Eastman film
1885

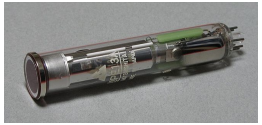

vidicon tube

CCD
sensor

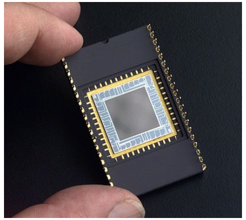

Nikon QV-1000C
Electronic Camera
1988

cellphone cameras

holography

X-ray CT

electron microscopy

ptychography

neutron imaging

EHT Telescope

Webb Telescope

Chemical Imaging

Electronic Imaging

Computational Imaging

*Reproductions from Wikipedia

What is Computational Imaging?

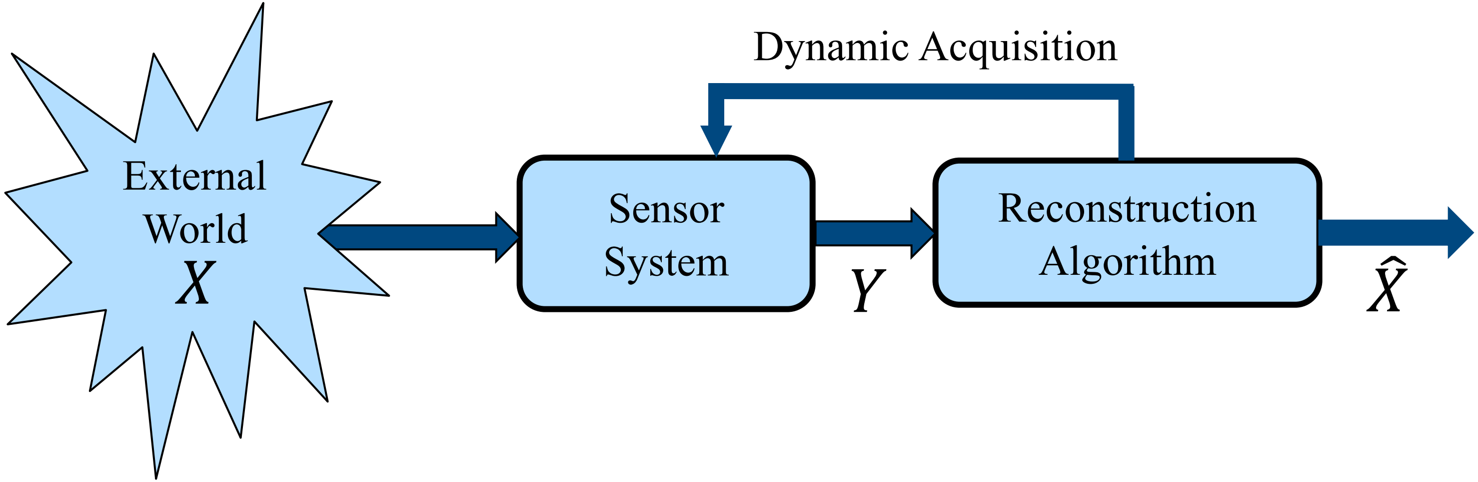

- Computational Imaging:
  - Engineering and science of turning data into images
  - Requires the solution to an inverse problem
- Co-design: Jointly optimize the sensor and algorithm
- Dynamic acquisition: Use results to optimize sensor parameters

Philosophy of Computational Imaging

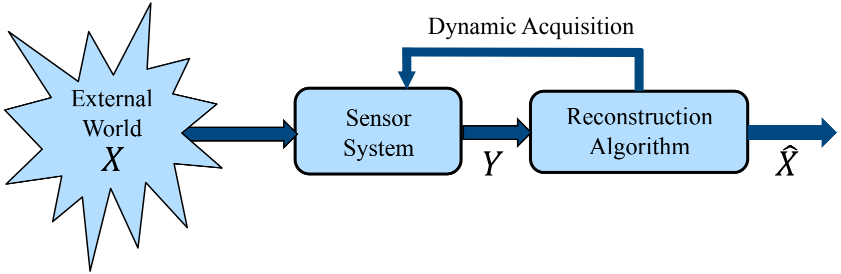

- Philosophy:
  - Traditional sensor design is reaching its limits
  - Make the **most informative** rather than **"purest"** measurement
- Mick Jagger's Theorem:
  - You can't always get what you want, but if you try sometimes, you might get what you need.

MBIR Reconstruction (Model-Based Iterative Reconstruction)

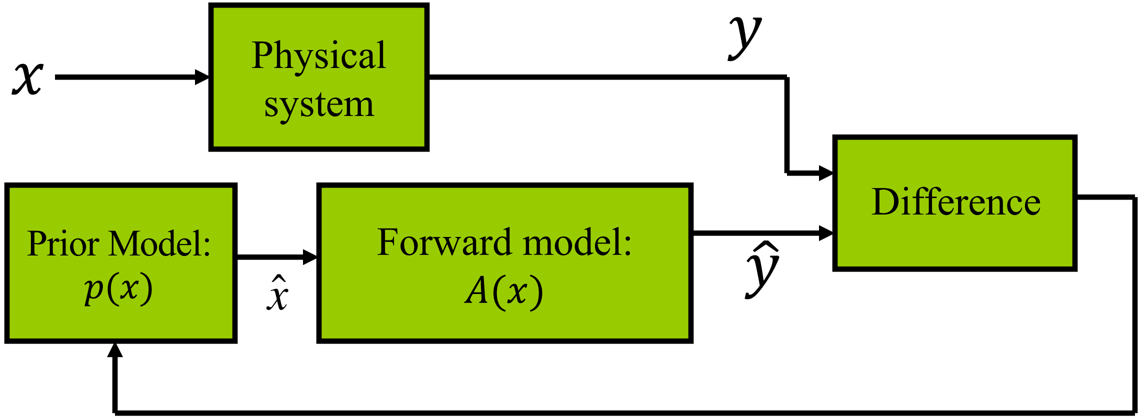

$$\hat{x} \leftarrow \arg\min_{x} \{\underbrace{\log p(y|x)}_{\text{forward model}} + \underbrace{\log p(x)}_{\text{prior model}}\}$$

$\hat{x}$ – Reconstructed object

y – Measurements from physical system

MBIR: “Thin Manifold” View

Sensor manifold –
Uses physics model

Physics sensor
manifold
$f(x) = -\log p(y|x)$

AI prior manifold
$h(x) = -\log p(x)$

Prior manifold –
Uses AI model

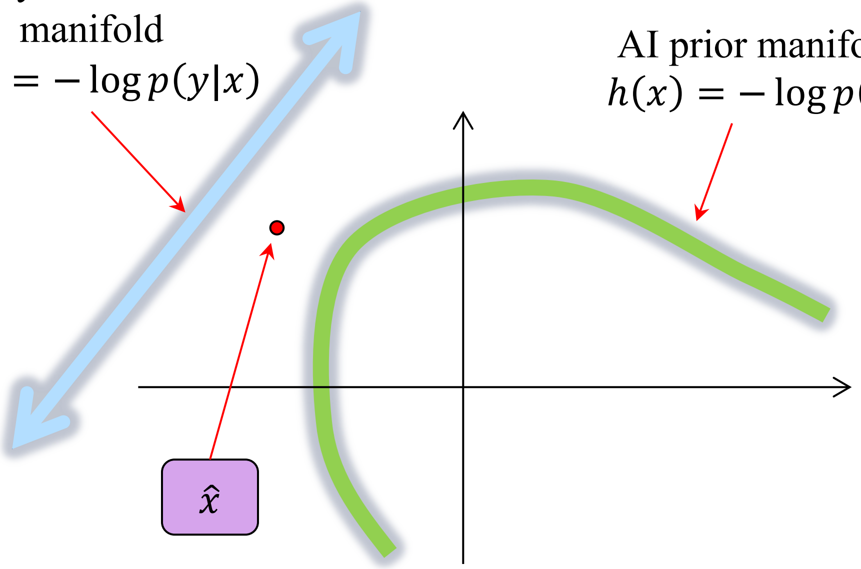

MBIR Reconstruction:

$$\hat{x} = \arg\min_{x}\{f(x) + h(x)\}$$

ML/AI in Computational Imaging

- Strengths/Weaknesses of ML/AI
- Direct versus Plug-and-Play Reconstruction
- The profound role of denoisers

Strengths/Weakness of ML/AI Methods

- Very detailed representations of data and systems
  - Can accurately model the distribution of real data
  - Can estimate the input/output behavior of complex systems
  - Can generate samples from sample distributions
- Not flexible
  - Difficult to adapt
  - Difficult to incorporate physics
  - Require copious training data

Two Approaches to AI in Computational Imaging

- Direct AI reconstruction
  - Fast and can be accurate with sufficient training
  - Not flexible (tuning hell)
  - Does not incorporate physics
- Plug-and-Play reconstruction
  - Models prior with AI denoiser
  - Alternates AI denoising with physics model
  - Modular and flexible

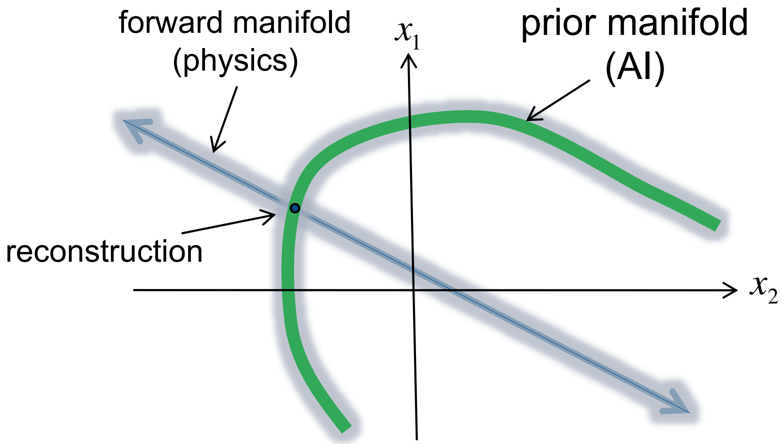

Variable Splitting

MBIR Reconstruction:

$$\hat{x} = \arg\min_{x=v} \{ f(x) + h(v) \}$$

- Split variables
- ADMM:
  - Alternating minimization of f and h
  - Augmented Lagrangian term enforces consistency.
  - Uses proximal maps

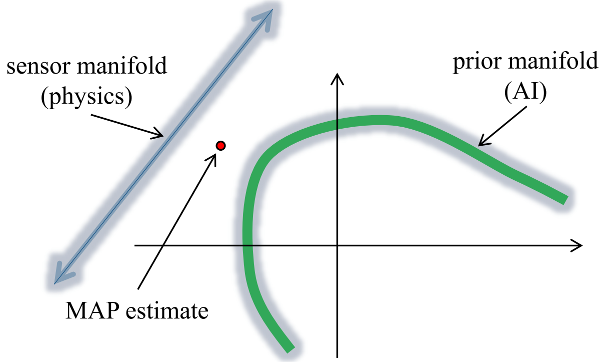

Proximal Maps

Minimize a function subject to a quadratic penalty on the distance (proximity) to a given base point.

- Proximal map of f with base point x:

$$\boxed{F(x) = \arg\min_{z} \left\{ \underbrace{f(z)}_{\text{Minimize a function}} + \underbrace{\frac{1}{2\sigma^2} \|z - x\|^2}_{\text{Quadratic distance penalty}} \right\}}$$

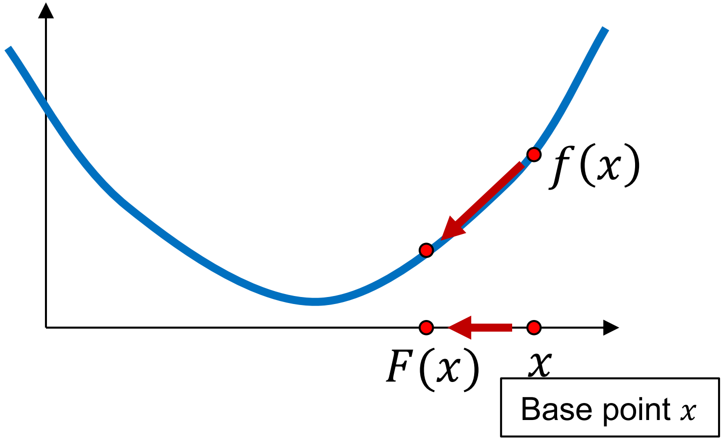

PnP Intuition and Convergence

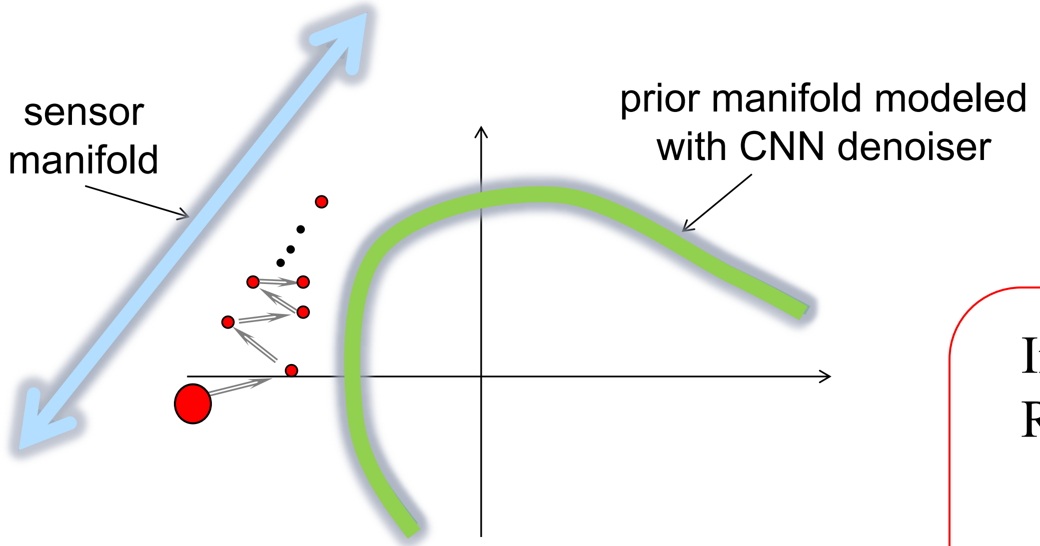

$$
\begin{aligned}
&\text{Initialize } v = x, u = 0 \\
&\text{Repeat } \{ \\
&\quad x \leftarrow F(v-u) \\
&\quad v \leftarrow \text{Denoise}(x+u) \\
&\quad u \leftarrow u + (x-v) \\
&\}
\end{aligned}
$$

[1] S. Venkat Venkatakrishanan, Charles A. Bouman, and Brendt Wohlberg, "Plug-and-Play Priors for Model Based Reconstruction," IEEE Global Conference on Signal and Information Processing (GlobalSIP), Austin, Texas, USA, December 3-5, 2013.

[2] Suhas Sreehari, S. Venkat Venkatakrishnan, Brendt Wohlberg, Gregery T. Buzzard, Lawrence F. Drummy, Jeffrey P. Simmons, and Charles A. Bouman, "Plug-and-Play Priors for Bright Field Electron Tomography and Sparse Interpolation," IEEE Transactions on Computational Imaging, vol. 2, no. 4, Dec. 2016.

PnP circa 2013

Forward model: match selected noisy subsamples

$$f(x) = \frac{1}{2}\|x - y\|^2$$

Ground Truth

Prior model: denoising algorithm

K-SVD

RMSE : 14.11

BM3D

RMSE : 12.56

Subsamples

Noise std. dev : 5% of max signal

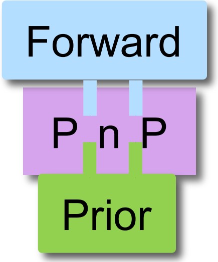

PLOW

RMSE : 14.54

TV

RMSE : 15.50

q-GGMRF

RMSE : 15.72

MBIR - Model Based Iterative Reconstruction

- Flow diagram
- Medical CT example

Computed Tomographic (CT) Imaging

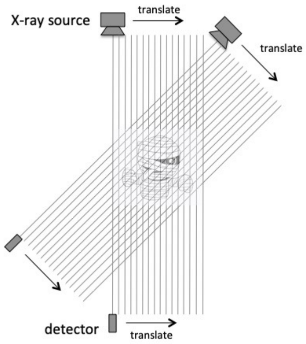

Parallel Beam CT: synchrotrons, electron microscopy, nano-X-ray sources

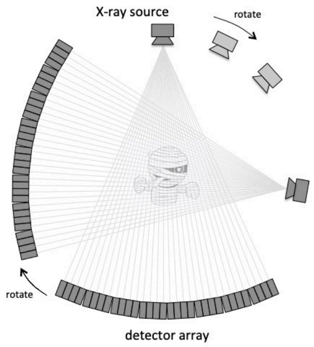

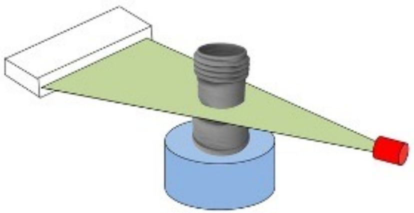

Fan Beam CT: Industrial CT

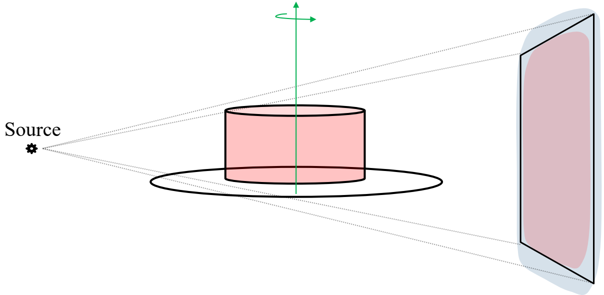

Cone Beam CT: Industrial CT, C-arm Scanners

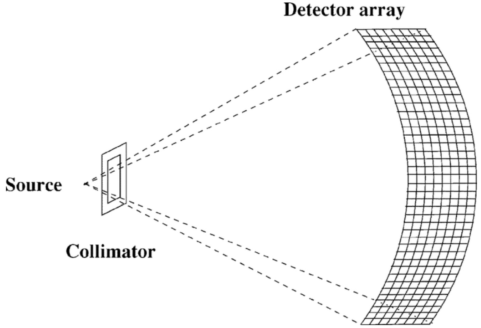

Multi-Slice Helical CT: Medical, transportation security

CT Forward Model

Volume to be Reconstructed

Noise

$$y = Ax + w$$

Measurements

System Matrix

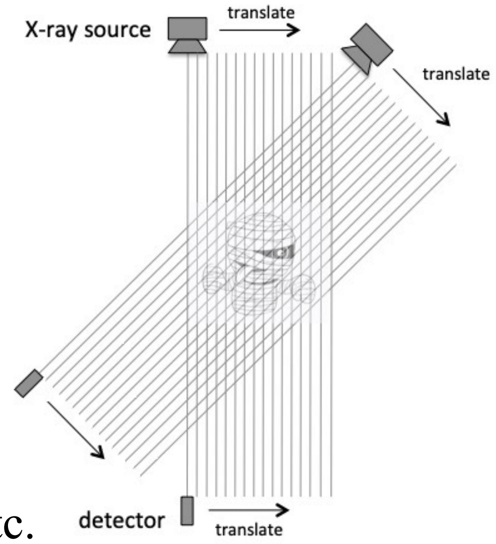

- **Problems:**
  - Not enough measurements: sparse or missing views, etc.
  - Low quality data: high noise, low dosage, short exposure, etc.
  - Model mismatch: metal, beam-hardening, scatter, poly-energetic, etc.
  - Resolution loss: detector blur, motion blur, X-ray spot size, etc.
- **Applications:**
  - Medical, scientific, industrial, and security

Model-Based Iterative Reconstruction (MBIR)

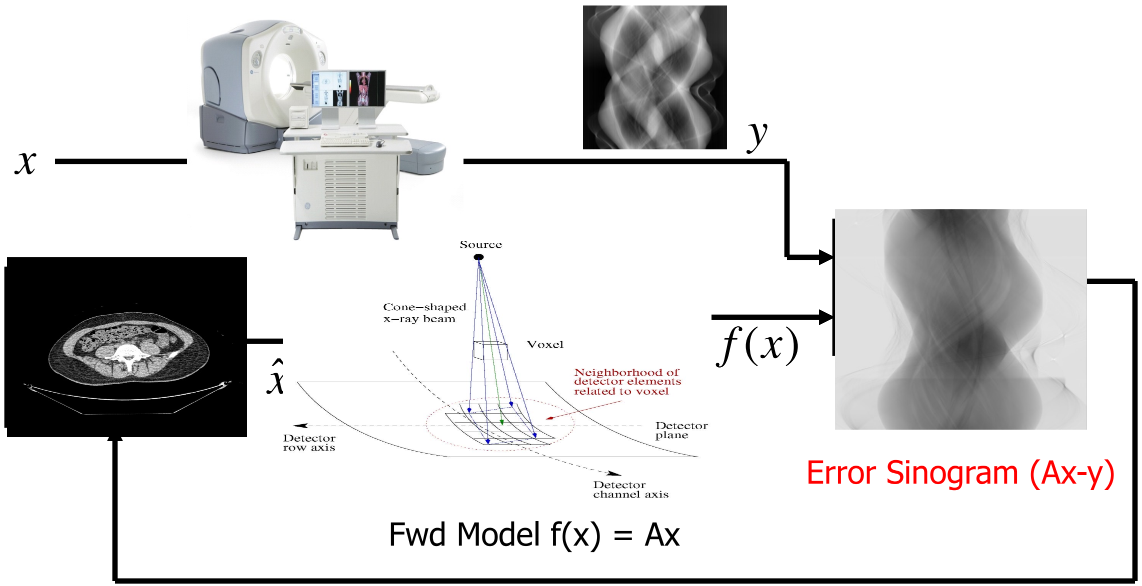

$$\hat{x} = \arg\min_{x \geq 0} \left\{ -\log p(y \mid x) - \log p(x) \right\}$$

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

MBIR for 64 slice GE VCT Scanner circa 2011

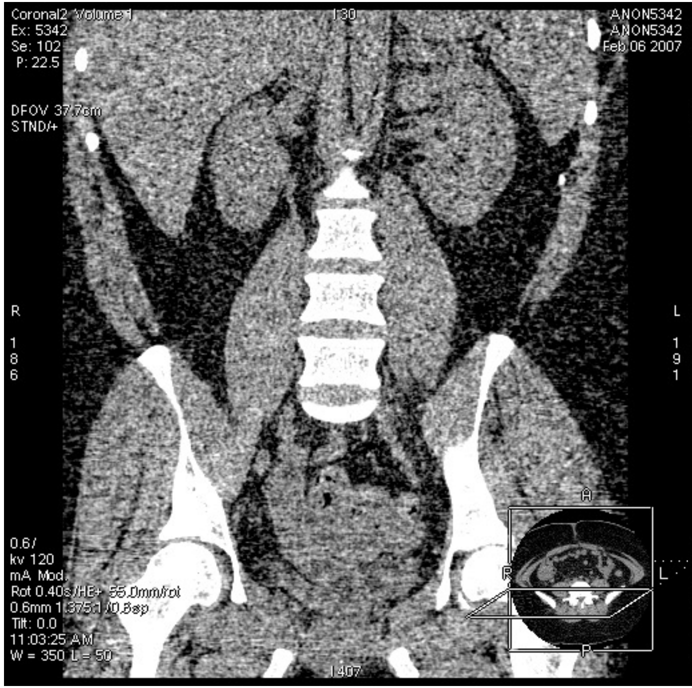

State-of-the-art 3D Recon

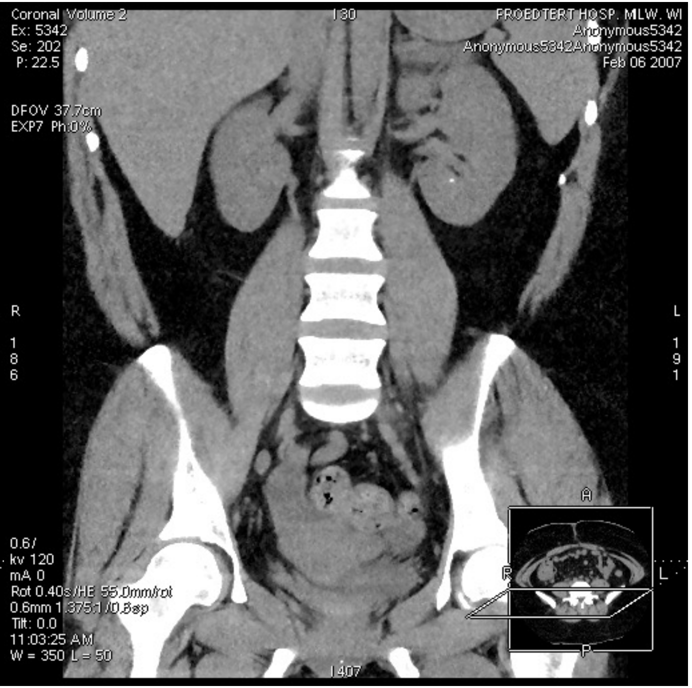

GE MBIR
Purdue/Notre Dame/GE algorithm

- Limitations:
  - Simple prior model
  - Very slow
  - Difficult to implement and use in scientific applications

Time Interlaced Model Based Iterative Reconstruction (TIMBIR)

K. Aditya Mohan, LLNL

John Gibbs, Northwestern University
Peter Voorhees, Northwestern University
Marc De Graef, CMU
Xianghui Xiao, APS
Charles Bouman, Purdue

Synchrotron Imaging of Time-Varying Sample

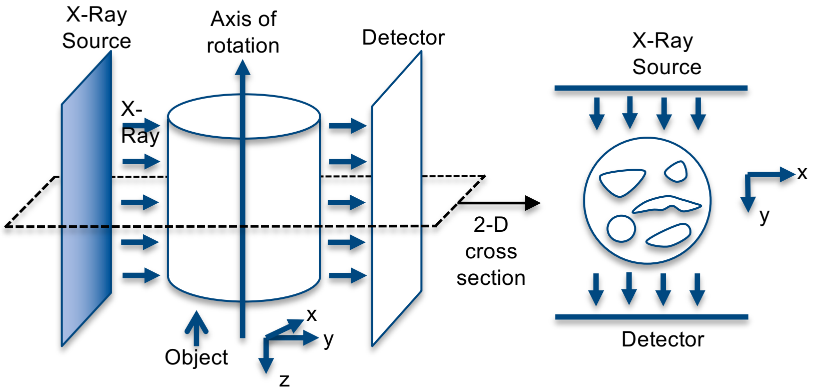

Real Synchrotron Projection Data

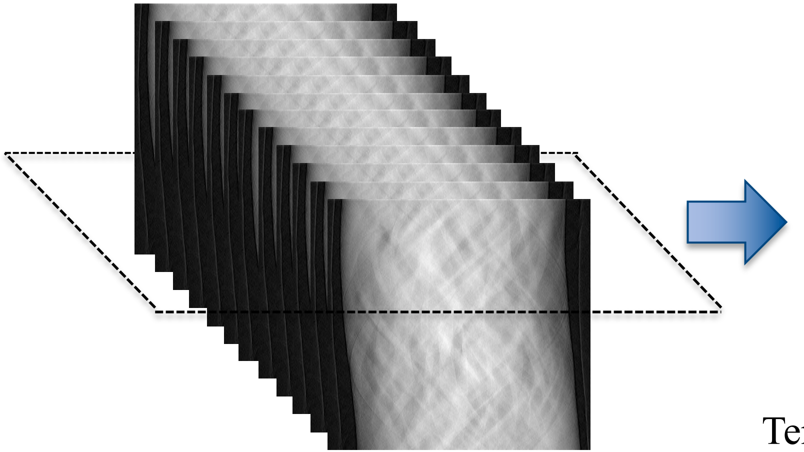

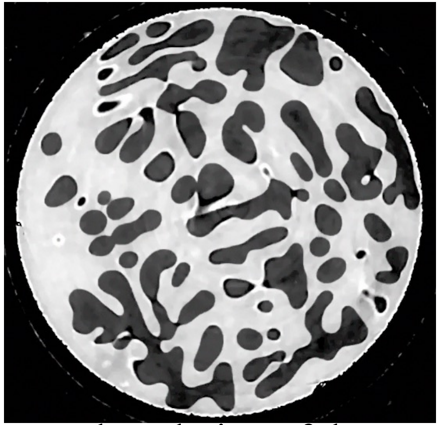

Temporal evolution of the sample

Examples of Interlaced Views

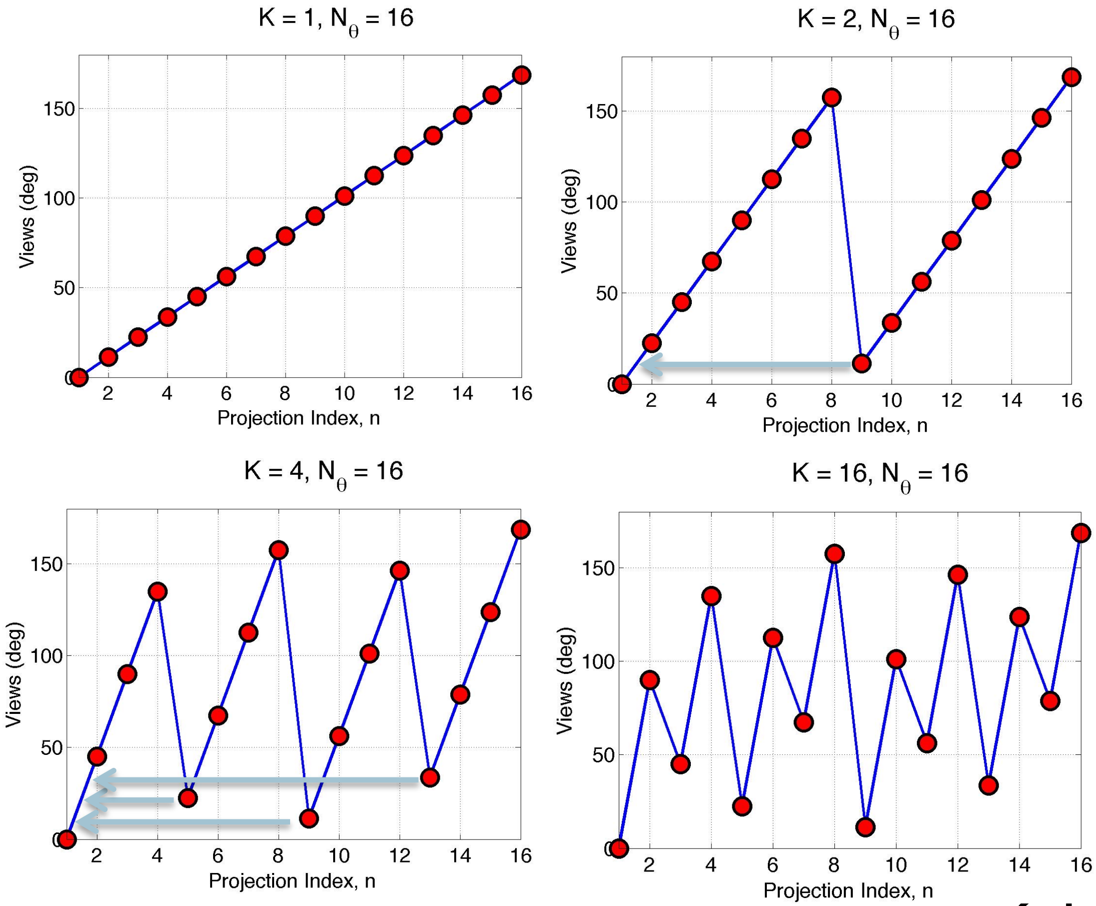

- Total number of discrete angles used is a constant.

  The time taken for rotation of object by 180 degrees decreases as K increases (or L decreases).

4D Prior Model in Space and Time

- Markov Random Field based model

$$-\log p(x) = \sum_{j=1}^{L} \sum_{\{k,l\}\in N} w_{kl}\rho\left(x_{j,k} - x_{j,l}\right) + \sum_{k=1}^{P} \sum_{\{j,i\}\in \Gamma} w_{kl}\rho\left(x_{j,k} - x_{i,k}\right) + \text{constant}$$

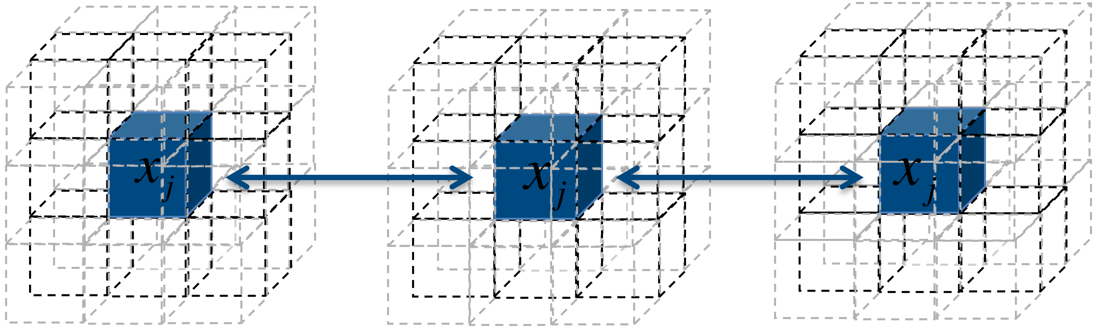

26 spatial neighbors
and 2 temporal neighbors of x_j

TIMBIR Experimental Results:

- 16x Speed-Up

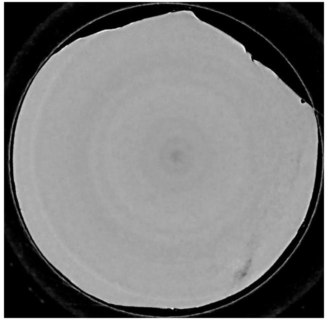

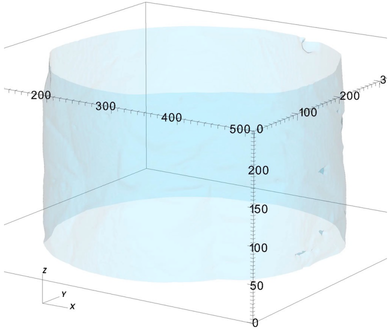

APS Experiment

- Solidification of aluminum and copper mixture
- Temperature decreased at 2^0 Celsius per minute
- 2000 views in a frame, interlaced over 16 sub-frames
- 16x speed up

Reconstruction

- (2048 x 2048 x 1000) space x 16 time
- $(0.65\ \mu m)^3$ voxel size
- 1.8 sec time step
- Image scaling: 10000 HU to 60000 HU

- 32x Speed-Up (k=16)

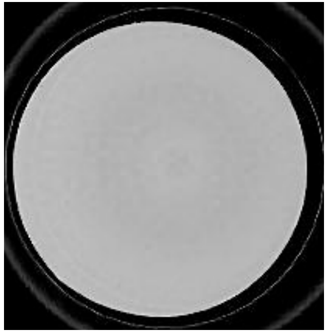

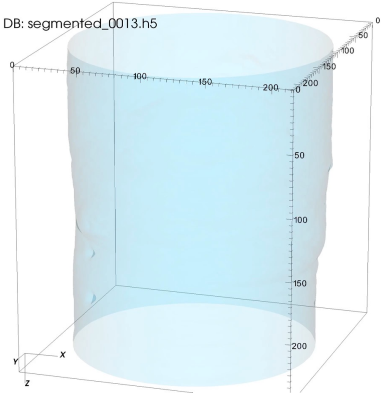

APS Experiment

- Solidification of aluminum and copper mixture
- Temperature decreased at 5^0 Celsius per minute
- 2000 views in a frame, interlaced over 32 sub-frames
- 32x speed up

Reconstruction

- (2048 x 2048 x 1000) space x 16 time
- $(0.65\ \mu m)^3$ voxel size
- 0.9 sec time step
- Image scaling: 10000 HU to 60000 HU

Neutron Imaging

Thilo Balke, Purdue/LANL

Alexander M. Long, LANL

Sven C. Vogel, LANL

Brendt Wohlberg, LANL

Charles A. Bouman, Purdue

Overview

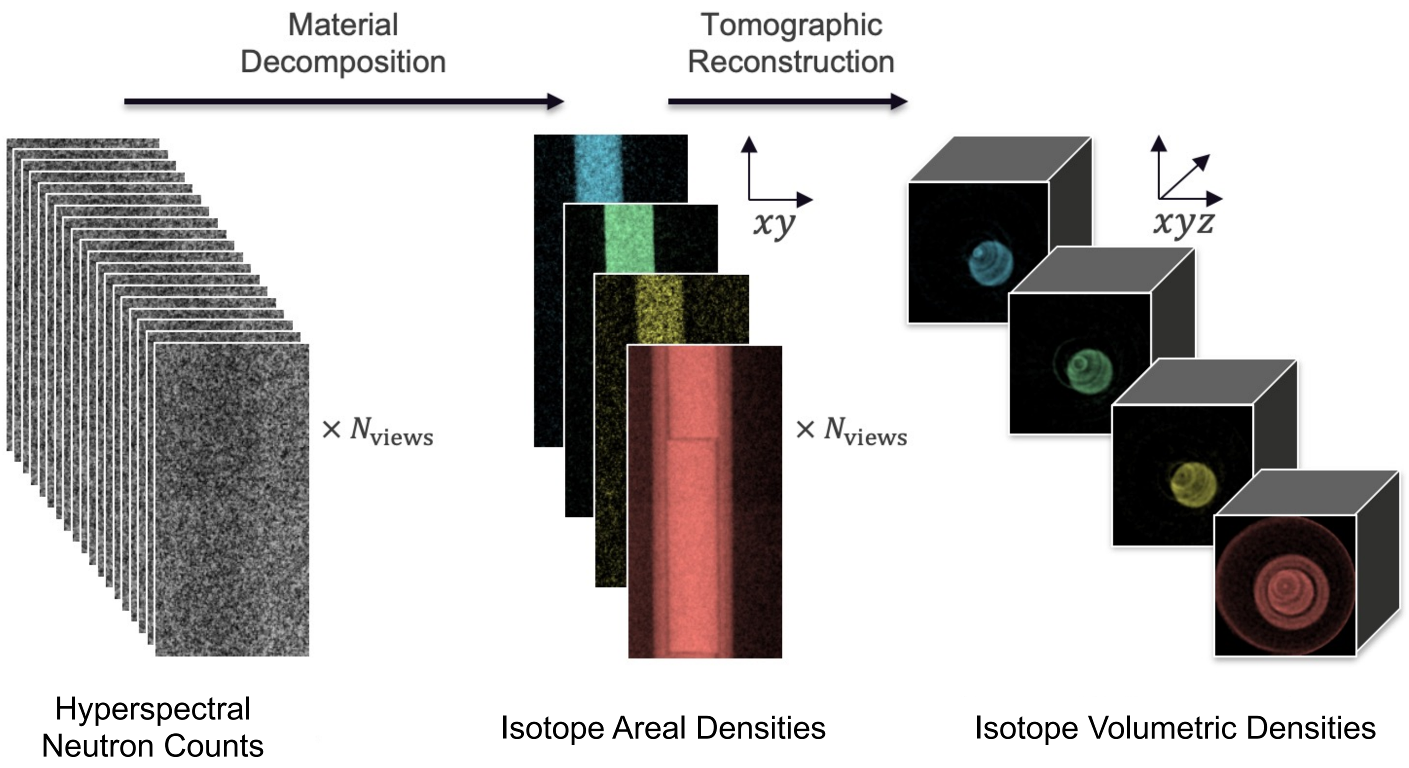

Material Decomposition
Tomographic Reconstruction
xy
xyz
$\times N_{\text{views}}$
$\times N_{\text{views}}$
Hyperspectral Neutron Counts
Isotope Areal Densities
Isotope Volumetric Densities

Measurement Model

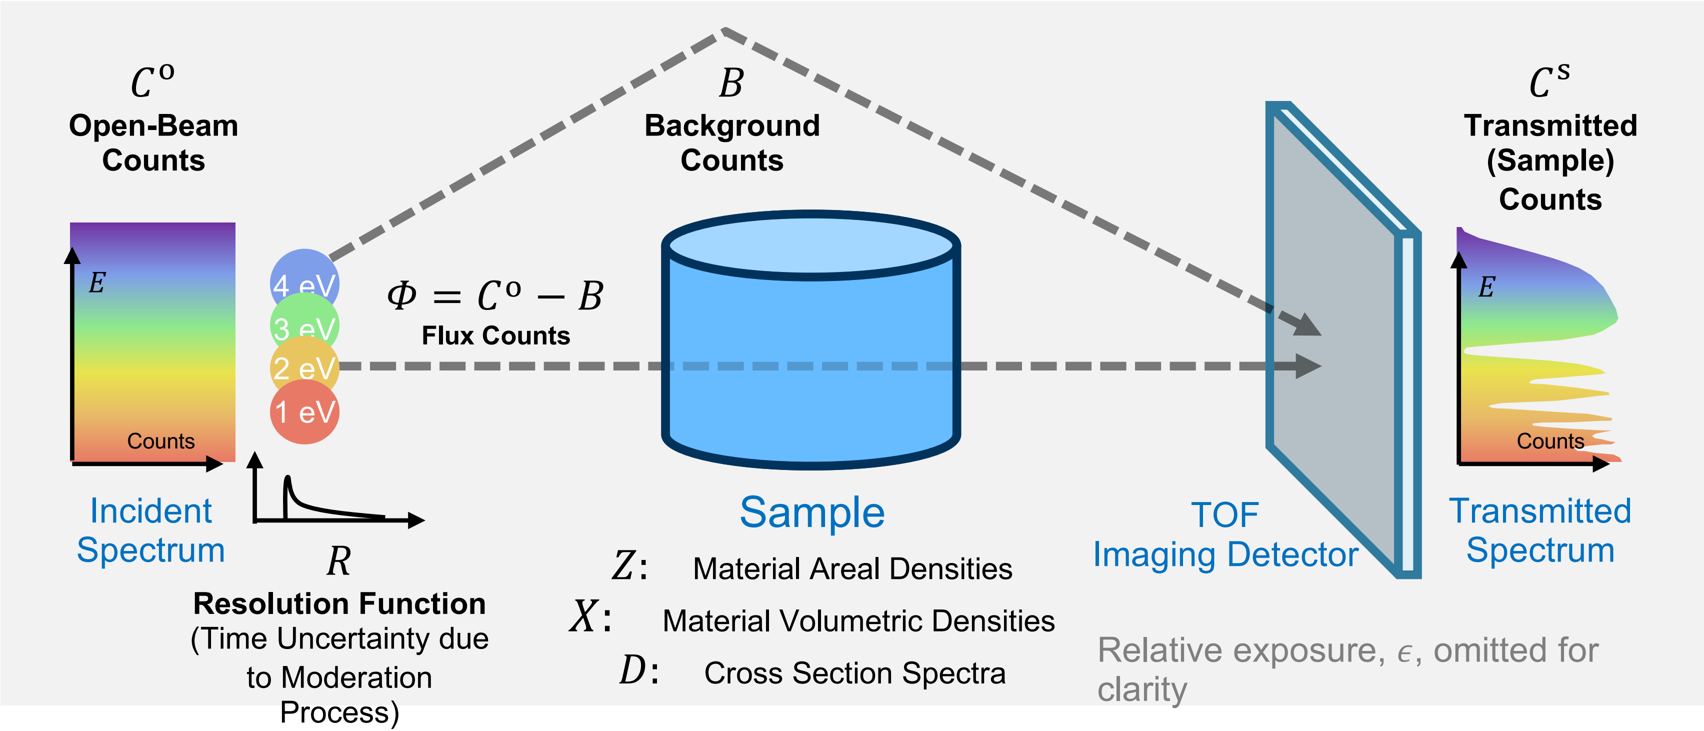

- Detector collects 2000 frames at 30nsec per frame (33 MHz) designed by Anton Tremsin, Space Sciences Laboratory, University of California, Berkeley

Forward Model

Dictionary of mass attenuation functions

Blur

Scatter

Tomographic Projection

Tomography:

$$Z = AX$$

$$\epsilon \rightarrow \mathrm{diag}(\boldsymbol{\epsilon})$$

Blank Scan

Object Scan

Forward Model:

$$C^{\mathrm{o}} = \Phi + B$$

$$C^{\mathrm{s}} = \epsilon[\Phi \odot (\exp(-ZD)R) + B]$$

$$\mathfrak{C}^{\mathrm{o}} \sim \mathrm{Poisson}(C^{\mathrm{o}})$$

$$\mathfrak{C}^{\mathrm{s}} \sim \mathrm{Poisson}(C^{\mathrm{s}})$$

- Optimization performed using accelerated proximal gradient method (APGM) with robust line search as part of the scientific computational imaging code (SCICO) software package.

Tomographic Reconstruction

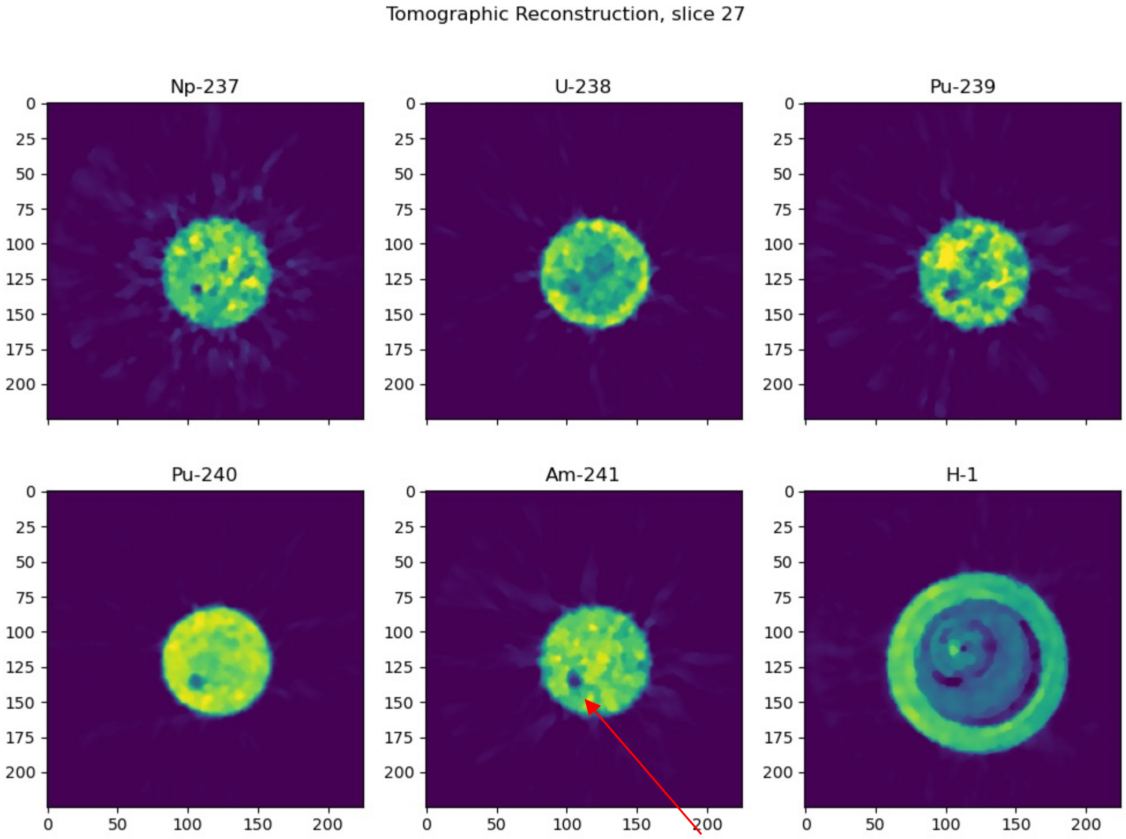

No cupping

No Halo artifacts

Void visible

Neutron Imaging

Diyu Yang, Purdue
Shimin Tang, LANL
Singanallur V. Venkatakrishnan, LANL
Mohammad S. N. Chowdhury, Purdue
Yuxuan Zhang, LANL
Hassina Z. Bilheux, LANL
Gregery T. Buzzard, Purdue
Charles A. Bouman, Purdue

Conventional View Selection in CT

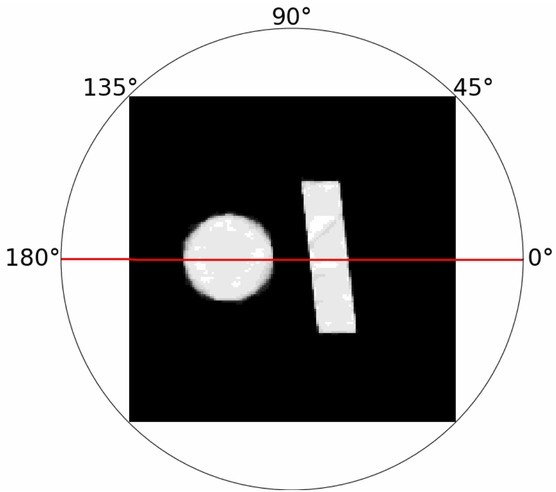

Equally spaced view angle

$$\Delta\theta = \frac{180^\circ}{N_v}$$

$$\theta_n = n \cdot \Delta\theta, n = 0, \ldots, N_v - 1$$

Good

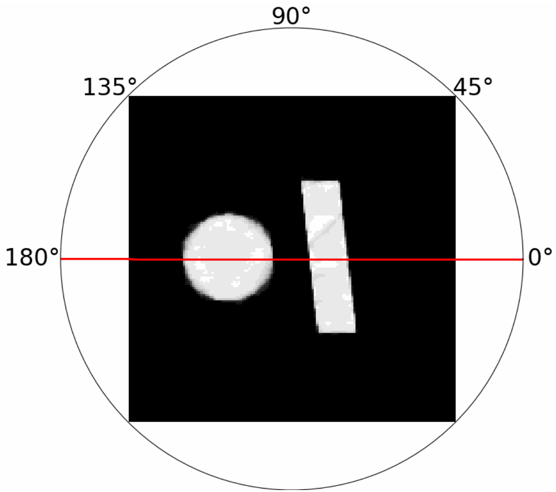

Golden angle

$$\theta_{golden} \approx 137.51^\circ$$

$$\theta_n = \left(n \cdot \theta_{golden}\right) mod\ 180$$

Better

Golden-Ratio

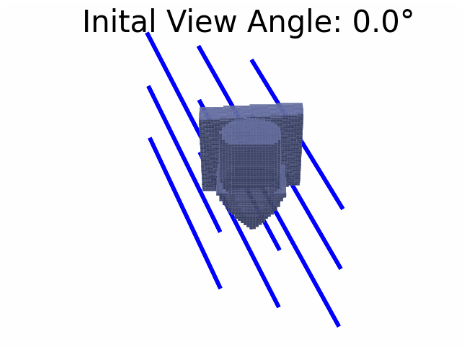

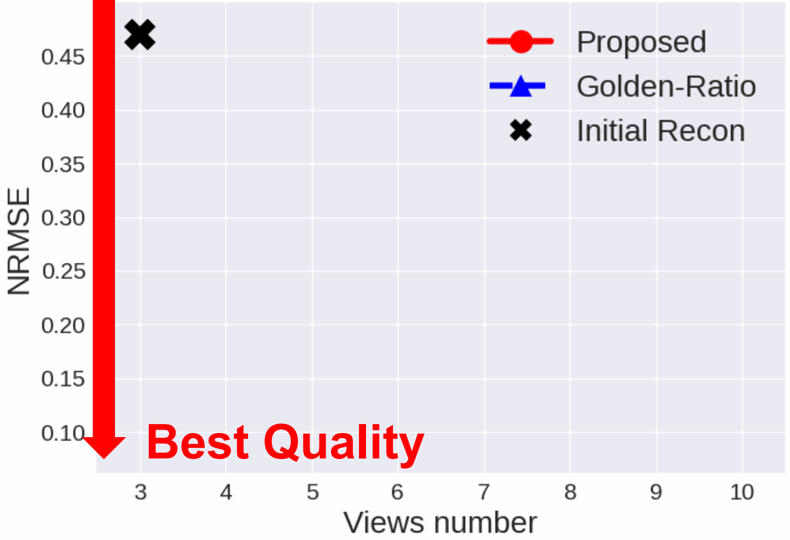

Dynamic

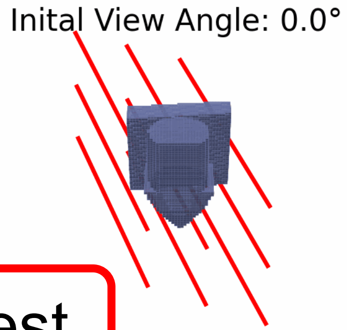

Best

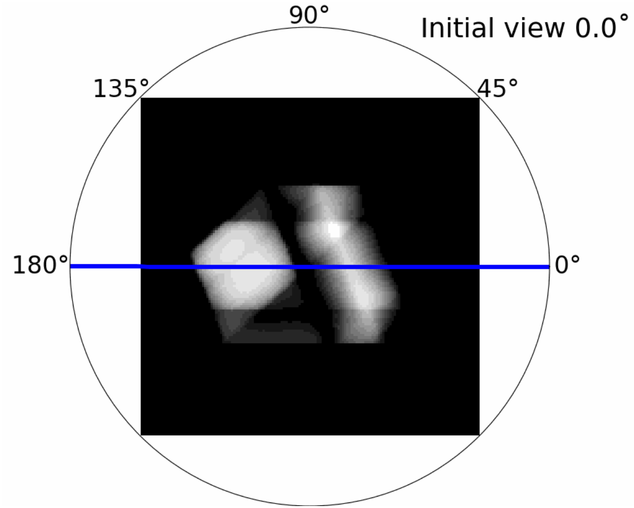

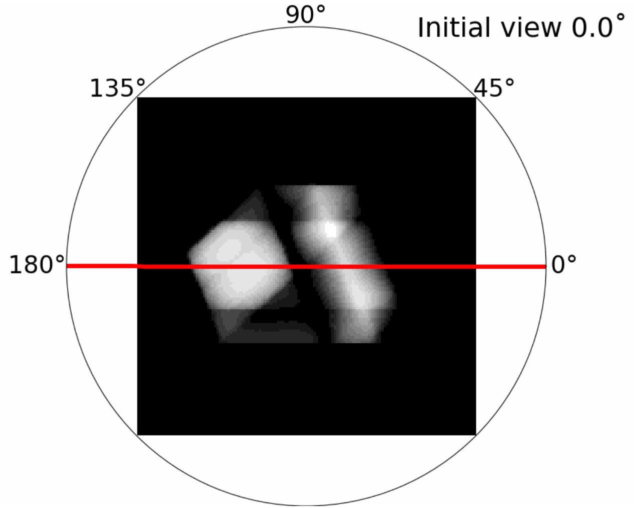

4D Recon using PnP/MACE

Soumendu Majee, Purdue
Thilo Balke, Purdue
Craig A. J. Kemp, Eli Lilly
Gregery T. Buzzard, Purdue
Charles A. Bouman, Purdue

4D MBIR Reconstruction

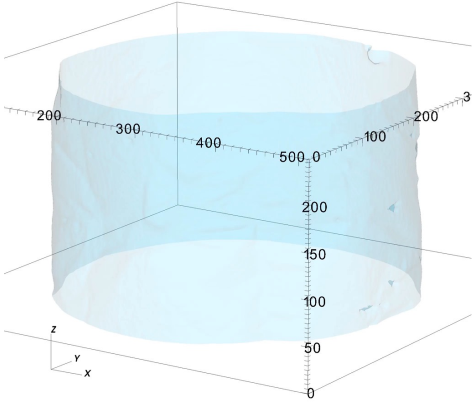

TIMBIR:
- Showed 16x increase in temporal resolution
- Based on simple 4D MRF prior

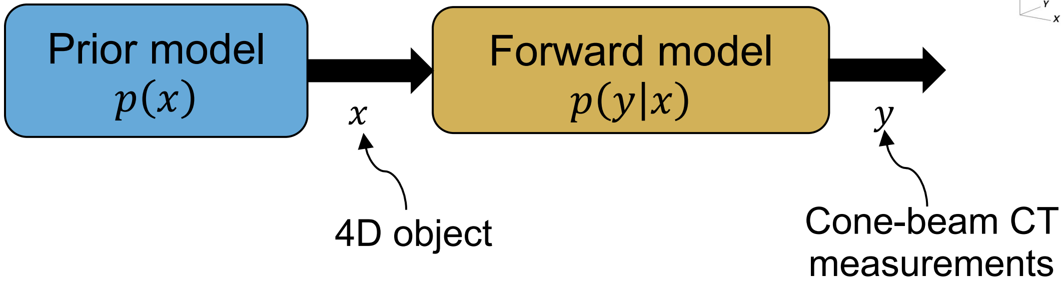

4D MBIR reconstruction:

$$\hat{x} \leftarrow \arg\min_{x}\{-\log p(y|x) - \log p(x)\}$$

Can we do better with advanced 4D priors?

Multi-Slice Fusion

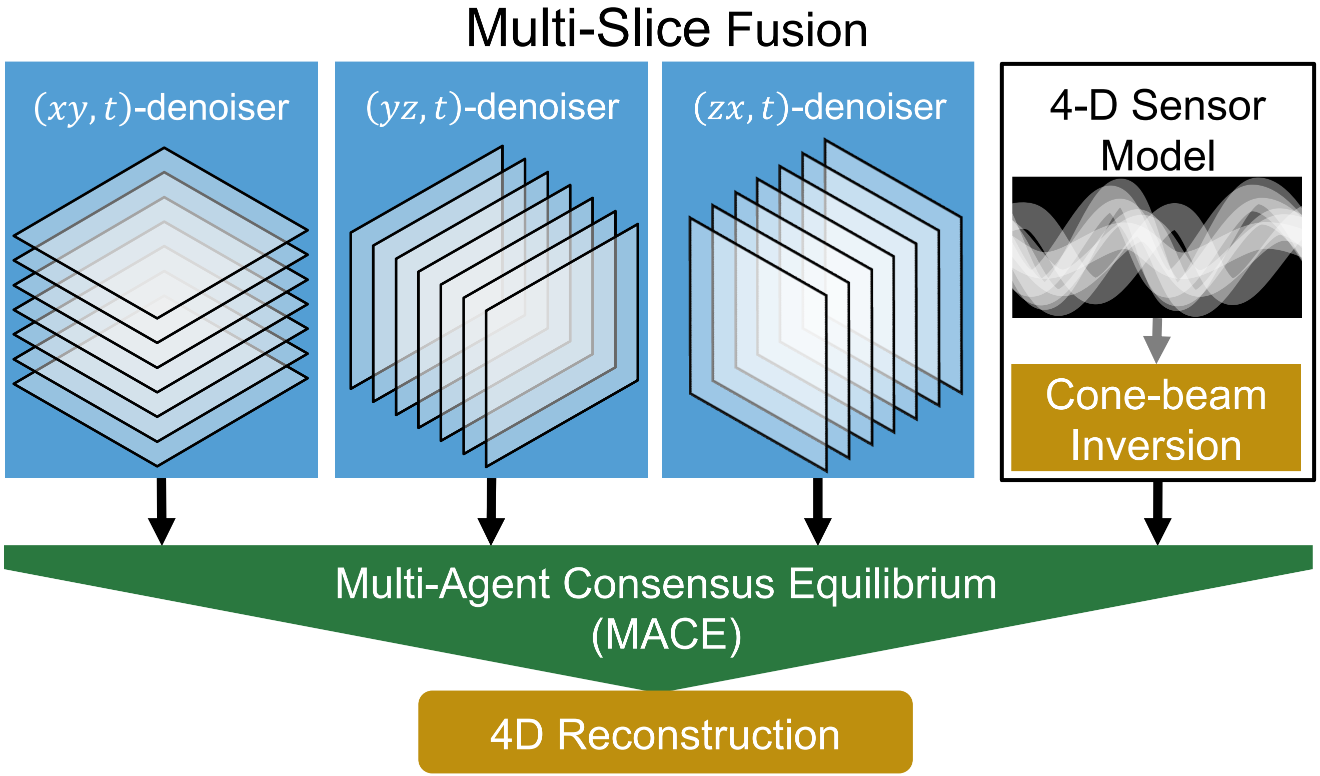

- Fuse multiple CNN denoisers to implement 4D prior
- Use 2D convolutions: fast and implementable
- No 4D training data required

Experimental Setup

Scanner Model	North Star Imaging X50
Source-Detector Distance	839 mm
Magnification	5.57
Cropped Detector Array	$731{\times}91$, $(0.254\ \text{mm})^2$
Detector resolution at ISO	45.7 μm
Number of Views per Rotation	150
Voxel Size	$(45.7\ \mu\text{m})^3$
Reconstruction Size (x, y, z, t)	$731{\times}731{\times}91{\times}16$

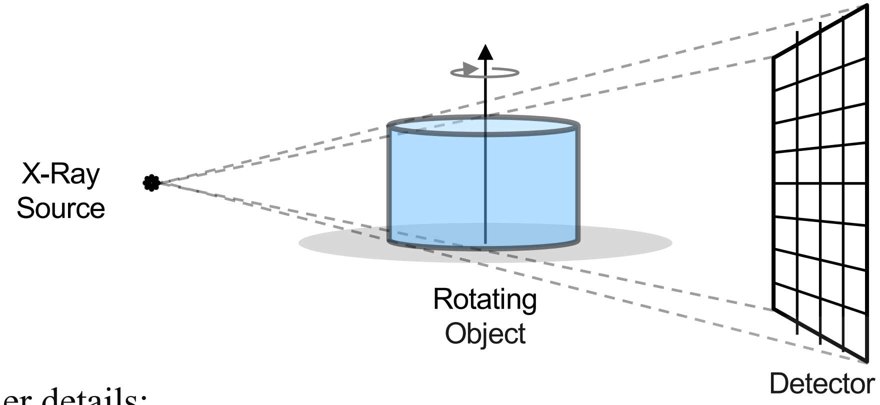

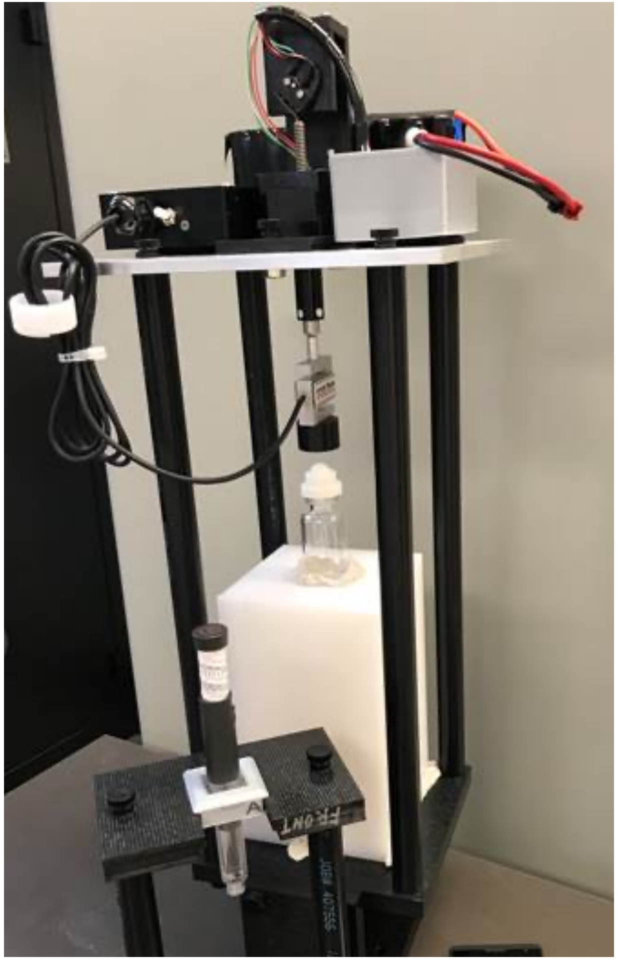

Other details:

- Object held in place by fixtures: artifacts
- All 4D results undergo preprocessing to correct for jig artifacts

Multi-Slice Fusion: Qualitative Comparison

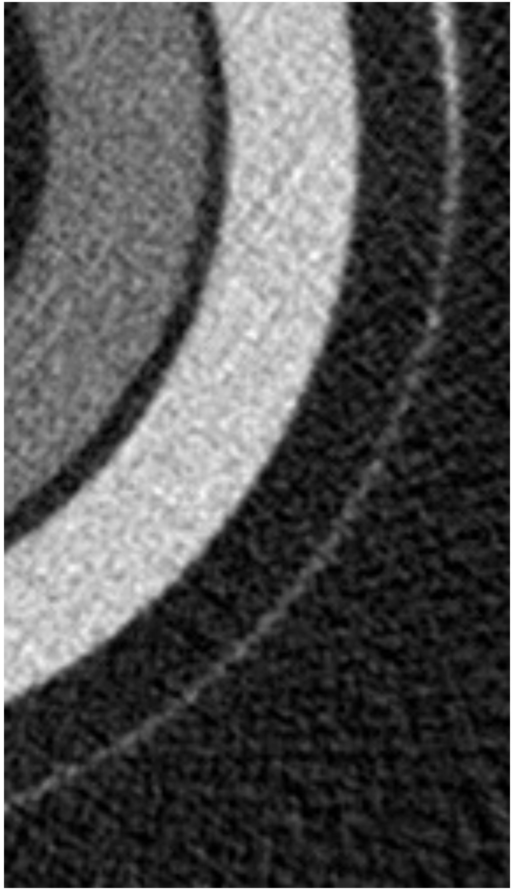

FBP (3D)

MBIR with 4D prior

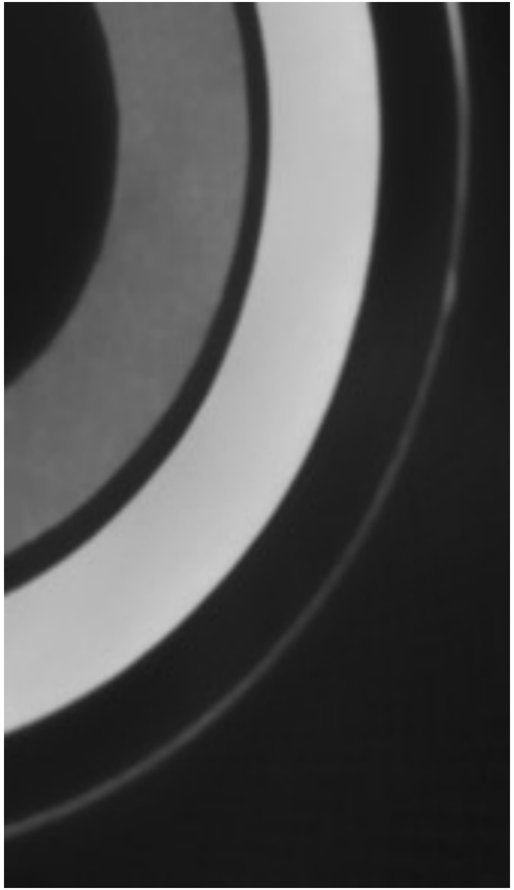

Multi-Slice Fusion

Vial Scan with Force-Curve

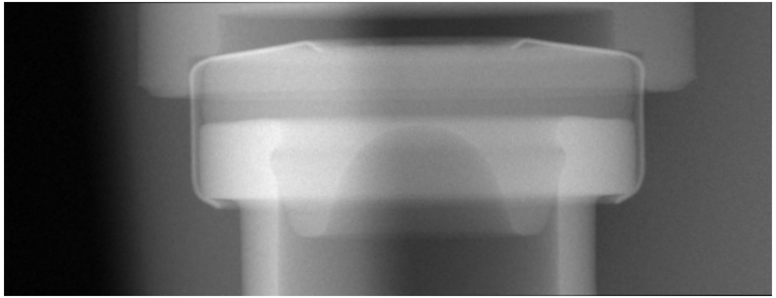

Sinogram View

- Scanner parameters:
  - 758 ×290 pixels, 3750 views, 25 full rotations
  - Detector spacing: 0.254×0.254 mm^2
  - Source-object distance: 152 mm
  - Object-detector distance: 695 mm
  - Magnification: ≈ 5.57
- Image Parameters (ROR)(rotations 5-8):
  - 758×758×290×4 voxels
  - Voxel size: $(0.05\ mm)^3$
  - Field of view: 38 mm (758 voxels)

Reconstruction (180° per time-point)

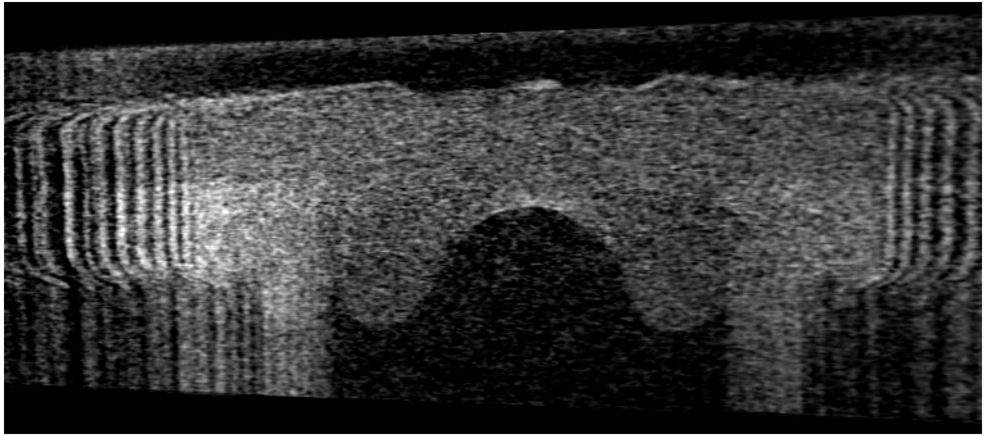

FBP

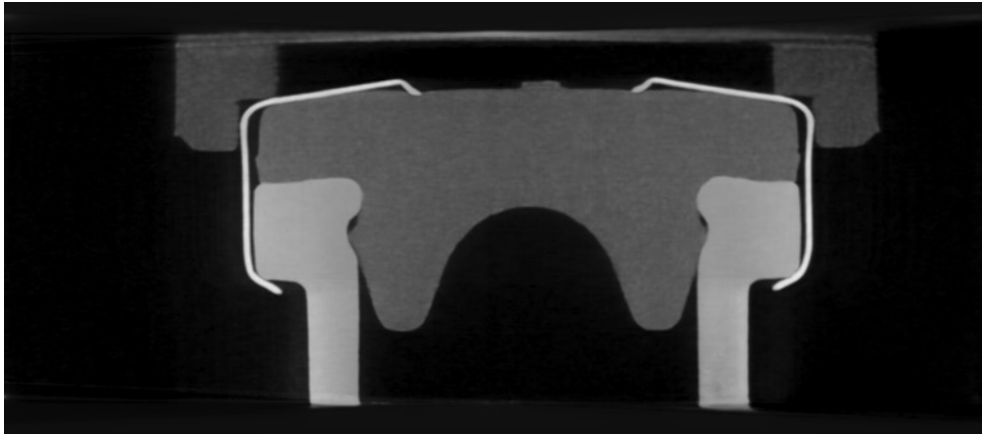

Multi-Slice Fusion

Digital Holographic Imaging and Wavefront Sensing

Casey Pellizzari, USAFA
Mark Spenser, AFRL
Charles A. Bouman, Purdue

Digital Holography

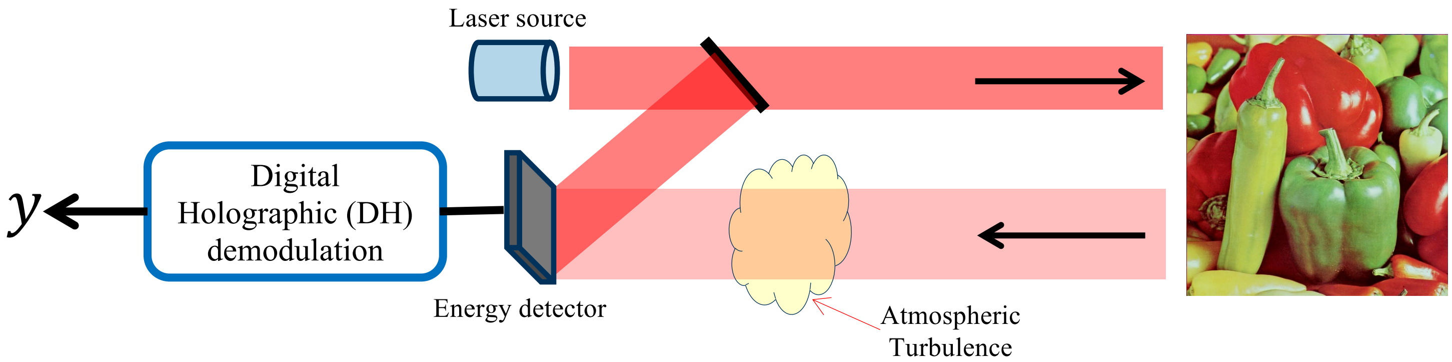

Sensing model:

$$\boldsymbol{y} = \boldsymbol{A_\phi} \boldsymbol{g} + \boldsymbol{w}$$

assumes Fourier demodulation

contains speckle

$\boldsymbol{y}$ – Complex measurement
$\boldsymbol{g}$ – Complex reflectance coefficient
$\boldsymbol{w}$ – Complex noise
$\boldsymbol{A_\phi}$ – Linear propagation model
$\boldsymbol{\phi}$ – Unknown phase distortion

DH Wavefront Correction

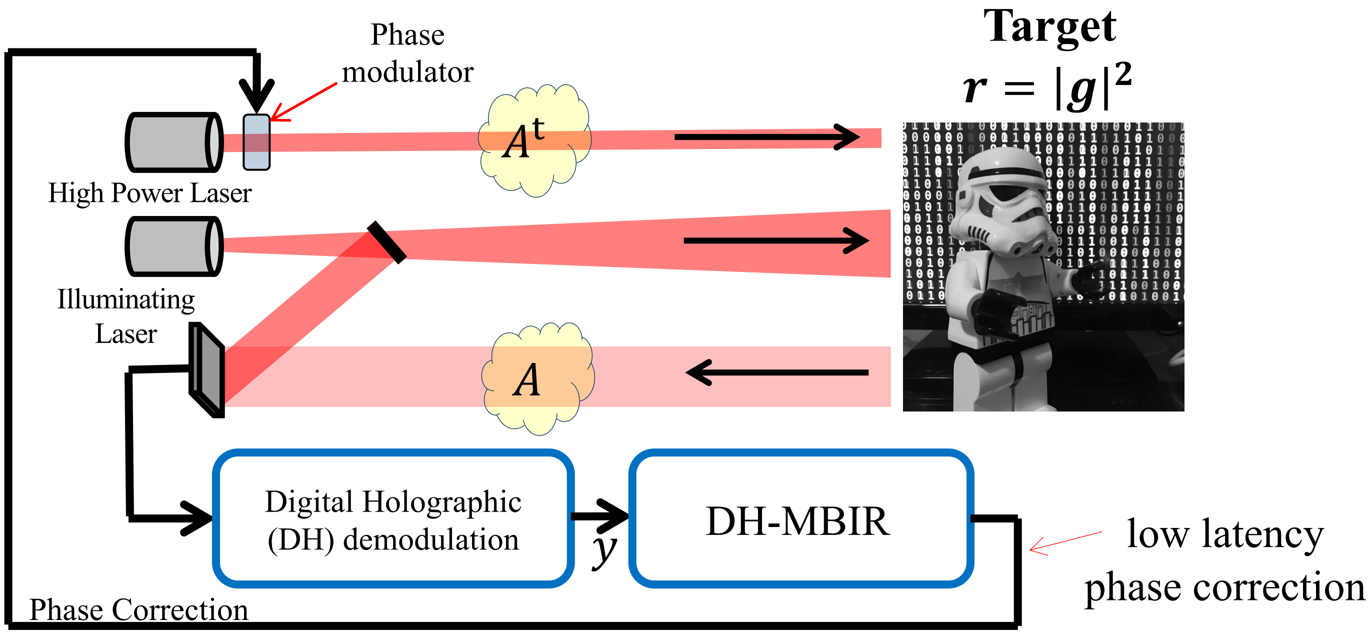

- Scary fact:
  - You can always increase throughput with more hardware
  - You <u>can't always reduce latency</u> with more hardware

DH-MBIR Reconstruction: Anisoplanatic Simulation Data

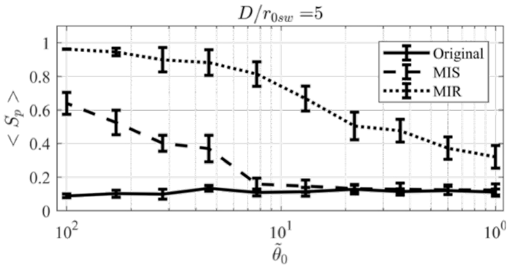

Peak Strehl ratio vs. $\tilde{\theta}_0$, averaged over 10 i.i.d. realizations

Simulation Parameters:

- 256x256 images
- 3 phase screens
- $\left(\frac{\lambda}{D}\right) = 1\, pixel$

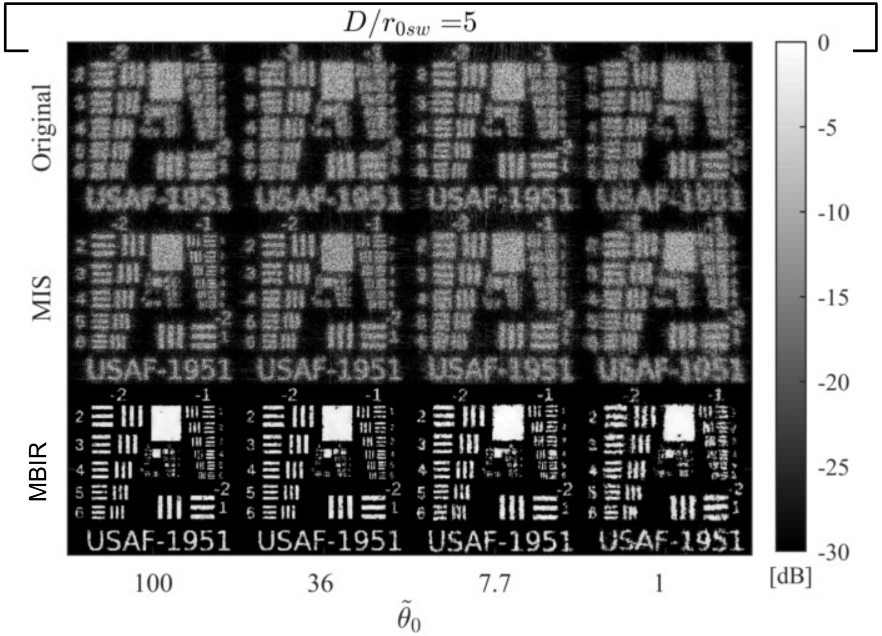

Experimental Reconstructions

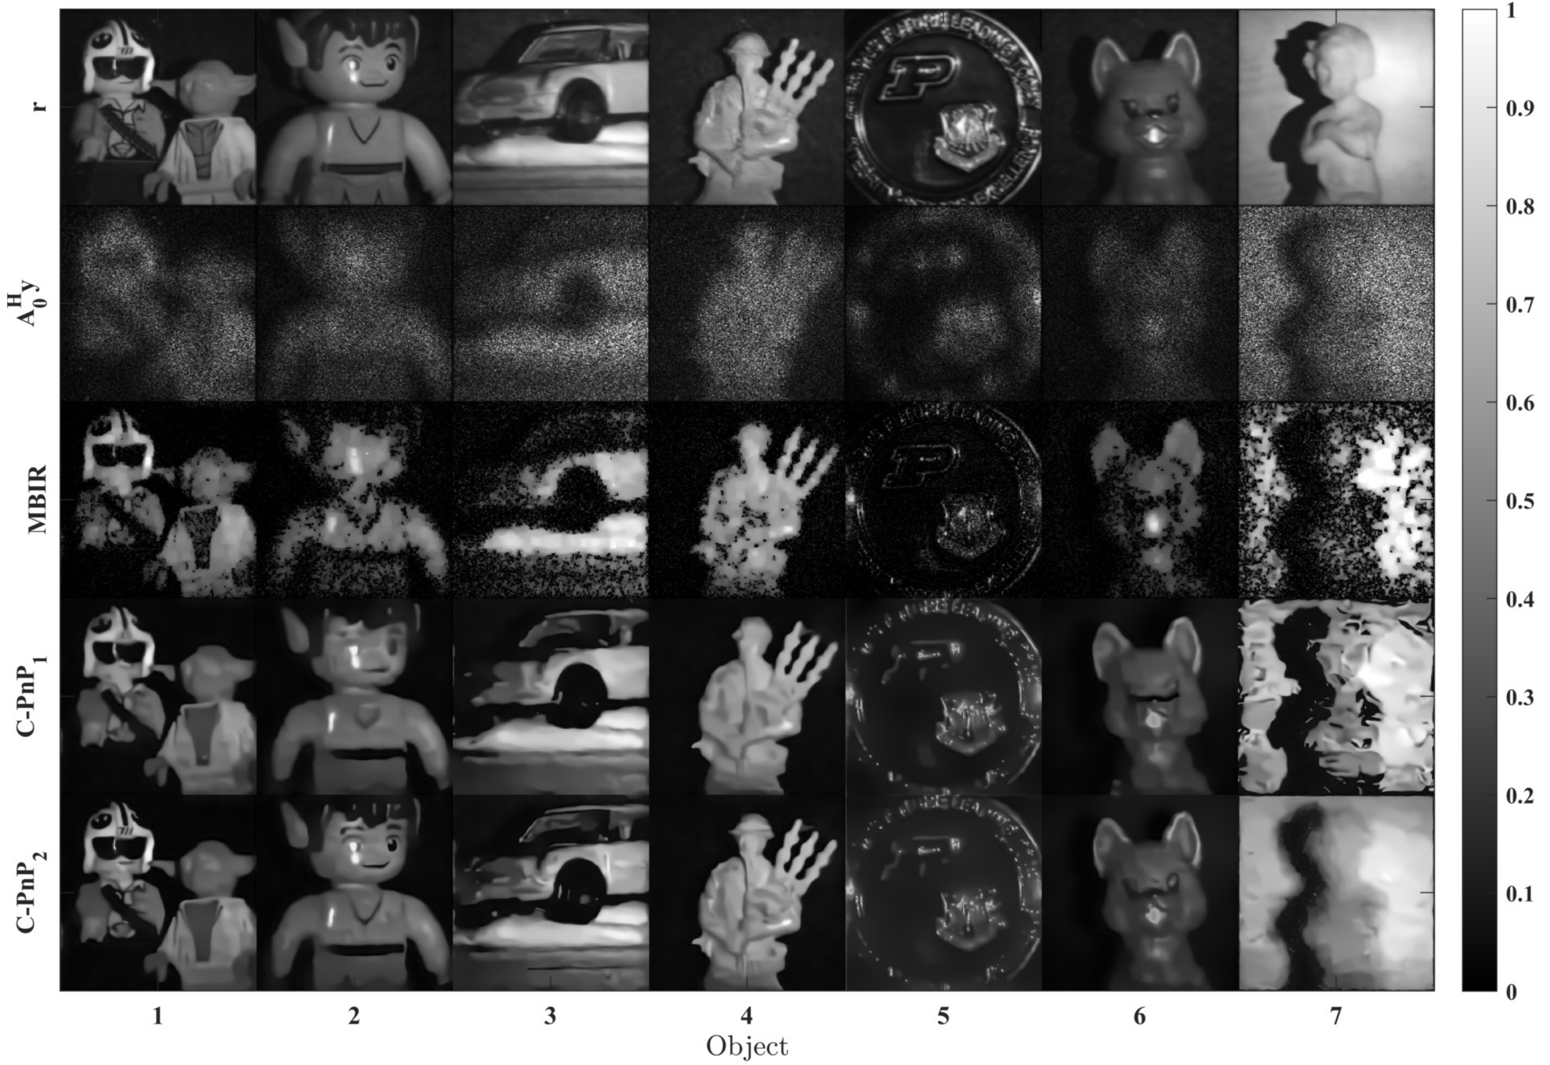

Take Aways

- Computational Imaging
  - Every sensing problem is an inverse problem
  - Make the most informative measurement
- Fusing physics-based and ML models
  - Plug-and-Play: Use denoiser as prior model
  - MACE: Integrate multiple physics and data models
- MBIR/PnP can be used to dramatically improve CT quality
  - Sparse data
  - Nonlinear forward problems
  - Dynamic measurements

What are the opportunities?

- How can co-design of algorithms/sensors improve performance?
- What opportunities exists for collaboration between our communities?

SV-MBIR Software Package

- Much (but not all) of this was done using SV-MBIR software
  - Open source BSD 3-clause software for:
  - Fast parallel and fan beam reconstruction of 3D volumes
    - Based on super-voxel algorithm that gives 100x to 1000x speedup
    - Nominated for 2017 Gordan Bell prize at SC17
  - Easy-to-use python interface
  - Fast multithreaded C code for multi-core CPUs
  - Available from conda-forge, PyPI or direct package installation.

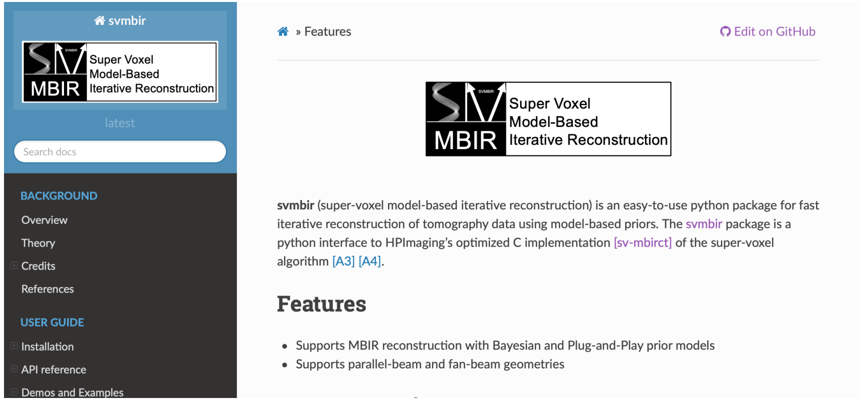

- URL:
  - https://github.com/cabouman/svmbir
  - https://svmbir.readthedocs.io/en/latest/install.html