Vectorized Coordinate Descent for Fast CT Reconstruction

Charles A. Bouman and Gregery T. Buzzard

Purdue University, School of ECE, BME, and Department of Mathematics

Main Message

Vectorized Coordinate Descent (VCD) is a fast, parallel algorithm for iterative CT reconstruction[1,2]. The VCD algorithm:

- Uses multi-granular voxel partitions to achieve fast convergence at all spatial scales.
- Has guaranteed convergence.
- Unifies the spectrum of methods from gradient descent to coordinate descent optimization.
- Is available via pip from MBIRJAX: an open-source package based on python jax for seamless use on CPUs and GPUs.

Each VCD iteration depends on the partition of the voxels into randomized subsets. If the voxels are all contained in a single subset, then VCD becomes Gradient Descent (GD) , and if each voxel is in its own subset VCD becomes Coordinate Descent (CD). Optimization is performed over a sequence of coarse and fine grain partitions.

Background

Sparse view tomography is widely used in scientific applications ranging from tilt beam microscopy [3] to neutron imaging [4]. Model-based iterative reconstruction (MBIR) is known to be among the best reconstruction algorithms, but its use has been limited by long reconstruction times.

Computed Tomographic (CT) Imaging:

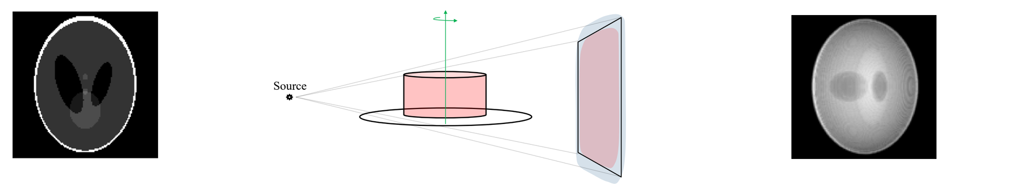

x: Physical object $\quad$ A: CT projection system $\quad$ y: Measured sinogram

$y = Ax + noise$: Measurement model

Model-Based Iterative Reconstruction (MBIR):

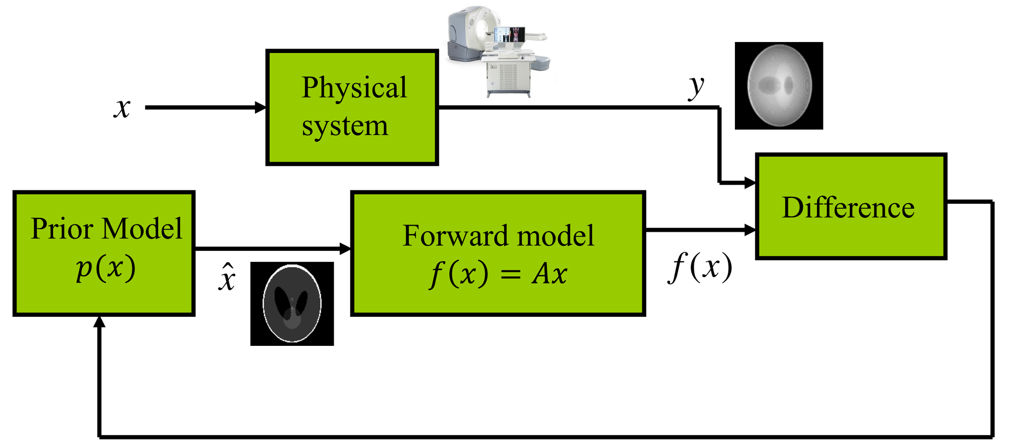

$$\hat{x} = \arg\min_x \left\{ \frac{1}{2} \|y - Ax\|_\Lambda^2 - \log p(x) \right\}$$

The Problem

Barriers to MBIR Adoption:

Too slow:
- 20 to 200x slower than FBP
- Large number of slow iterations

Too hard to use:
- Difficult to pick parameters
- Difficult to run software

Too difficult to implement:
- Need to "roll your own" for every geometry

An Algorithm that fits the Software and Hardware:

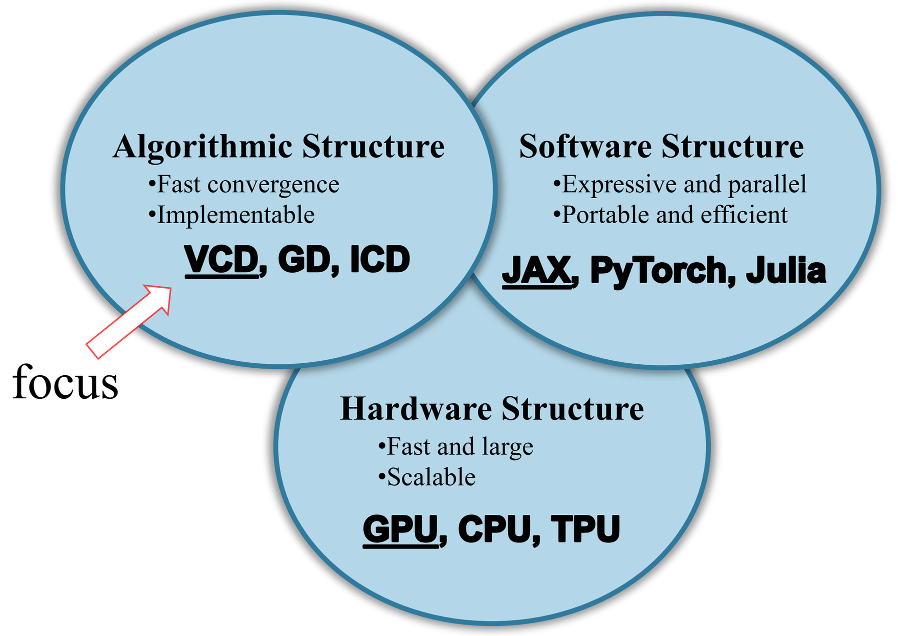

Algorithmic Challenge:

- Simplified problem:

$$\hat{x} = \arg\min_x \left\{ \frac{1}{2} \|y - Ax\|^2 \right\}$$

- Gradient Descent (GD):
  - ✗ Slow convergence
  - ✓ Fast updates
- Coordinate Descent (CD):
  - ✓ Fast convergence
  - ✗ Slow updates

Can we get the best of both worlds?

Vectorized Coordinate Descent

Coordinate Descent:

Update one pixel:

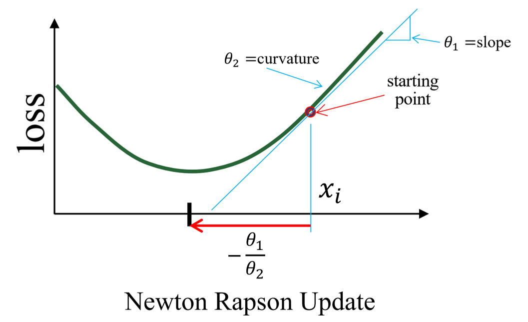

Newton Rapson Update

Not practical on GPU hardware!

```
# Define loss function
def: L(x) = 1/2 ||y − Ax||^2

# Initialize image
x ← 0

for each iteration:
    for each pixel i ∈ {0, ⋯, N − 1}:

        # Compute pixel update
        θ1 ← ∂/∂x_i L(x)
        θ2 ← ∂²/∂x_i² L(x)
        x_i ← x_i − θ1/θ2
Return x
```

Vectorized Coordinate Descent:

Partition pixels into subsets (colors)

For each color:

- Update pixels in parallel
- Use coordinate descent optimization

Cycle through subsets in sequence

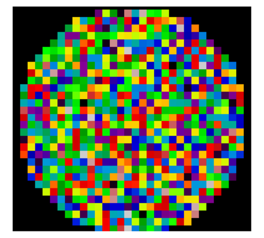

Bugzilla!

-Algorithmic instability that occurs when too many pixels are updated simultaneously.

-Conventional solution is to damp updates by a factor of α [5,6].

```
def: L(x) = 1/2 ||y − Ax||^2

# Initialize
(e, x) = (y, 0)
θ_{2,k} ← diag(A^t A)_{S_k}

for each iteration:
    for each subset k ∈ {0, ⋯, K − 1}:

        # Compute pixel updates
        θ1 ← −[A_k]^t e
        d ← −θ1 / θ_{2,k}
        x_{S_k} ← x_{S_k} + d

        #Update state
        e ← e − A_k d
Return x
```

Bugzilla Swatter:

```
for each subset k ∈ {0, ⋯, K − 1}:

    # Compute pixel updates
    θ1 ← −[A_k]^t e
    d ← −θ1 / θ_{2,k}
    x_{S_k} ← x_{S_k} + d

    #Update state
    e ← e − A_k d
```

Unstable ☹

Bugzilla Swatter

```
for each subset k ∈ {0, ⋯, K − 1}:

    # Compute image change
    θ1 ← −[A_k]^t e
    d ← −θ1 / θ_{2,k}

    # Compute optimal step
    p ← A_k d
    α* ← ⟨e, p⟩ / ⟨p, p⟩

    # Update state
    x_{S_k} ← x_{S_k} + α* d
    e ← e − α* p
```

Converges ☺

Unified Framework for GD and CD:

VCD always converges!

One algorithm for all:
- 1 subset ⇒ Gradient Descent
- N subsets ⇒ Coordinate Descent
- > 1 and $< N$ ⇒ Something new

Multi-Granular VCD

Multi-Granular VCD:

- Change the number of subsets during optimization
- Vary between GD and CD and everything in between
- A little like SGD, but the batches are in the domain

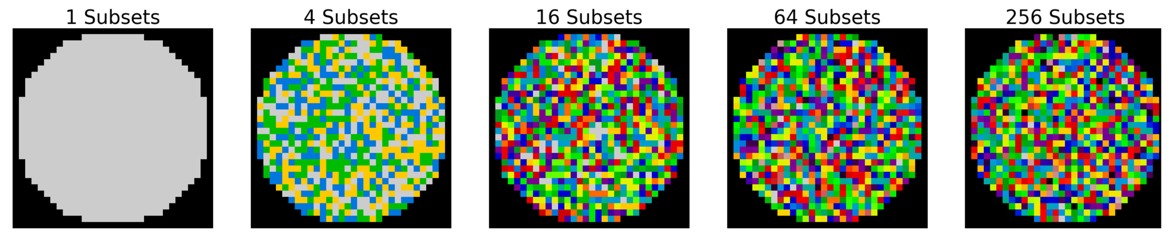

courser granularity
Gradient Descent
Better low frequency convergence

finer granularity
Coordinate Descent
Better high frequency convergence

MBIRJAX Software Package:

- Open-source BSD 3-clause
- Easy to use and install from PyPI
- Easy to create new geometries:
  - Parallel beam
  - Cone beam
  - More to come

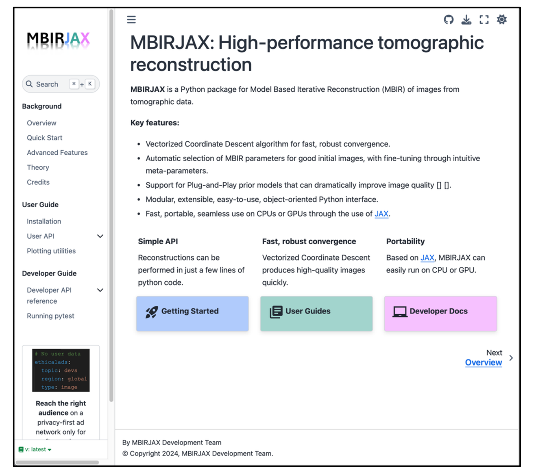

https://mbirjax.readthedocs.io
https://github.com/cabouman/mbirjax

Simulation Results:

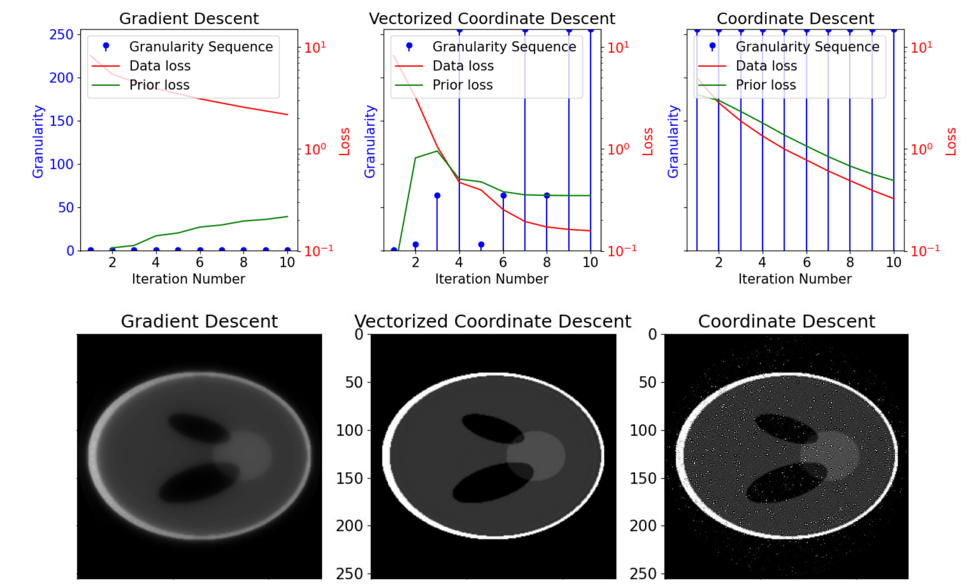

Experimental Results:

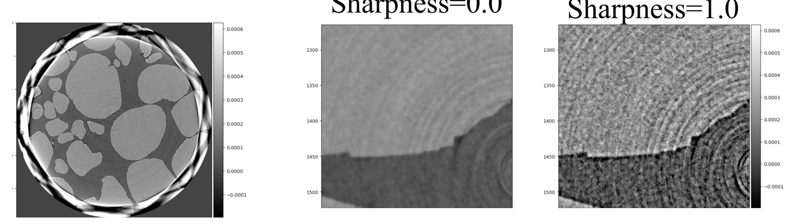

LBNL Advanced Light Source Reconstructions

Conclusions

Vectorized coordinate descent:
- Unifies Gradient and Coordinate Descent algorithms
- Updates pixels in parallel
- Guaranteed convergence (Bugzilla swatter)

Random pixel partitions:
- More subsets ⇒ better high frequency convergence
- Fewer subsets ⇒ better low frequency convergence

Multi-Granularity:
- Achieves good convergence at all frequencies
- Simpler to implement than multi-grid/multi-resolution

References

[1] Charles A. Bouman, *Foundations of Computational Imaging: A Model Based Approach*, SIAM 2022, Figure 8.4, page 119.

[2] Jean-Baptiste Thibault, Ken Sauer, Charles Bouman, and Jiang Hsieh, "A Three-Dimensional Statistical Approach to Improved Image Quality for Multi-Slice Helical CT," Medical Physics, pp. 4526-4544, vol. 34, no. 11, November 2007.

[3] S. Venkat V. Venkatakrishnan, Lawrence F. Drummy, Michael Jackson, Marc De Graef, Jeff Simmons, and Charles A. Bouman, "A Model Based Iterative Reconstruction Algorithm For High Angle Annular Dark Field - Scanning Transmission Electron Microscope (HAADF-STEM) Tomography," *IEEE Transactions on Image Processing*, pp. 4532-4544, vol. 22, no. 1, November 2013.

[4] Thilo Balke, Alexander M. Long, Sven C. Vogel, Brendt Wohlberg, and Charles A. Bouman, "TRINIDI: Time-of-Flight Resonance Imaging with Neutrons for Isotopic Density Inference," *IEEE Transactions on Computational Imaging*, vol. 10, pp. 154-169, 2024.

[5] J. A. Fessler, E. P. Ficaro, N. H. Clinthorne, and K. Lange, "Grouped-co-ordinate ascent algorithms for penalized-likelihood transmission image reconstruction," *IEEE Trans. Med. Imag.*, vol. 16, pp. 166–175, Apr. 1997.

[6] K. D. Sauer, S. Borman, and C. A. Bouman, "Parallel Computation of Sequential Pixel Updates in Statistical Tomographic Reconstruction," vol. 2, pp. 93-96, *IEEE Int'l Conf. on Image Proc.*, Washington, DC, Oct. 22-25, 1995.

Acknowledgements

Sponsors:
- Eli Lilly Company
- Oak Ridge National Laboratory
- The Showalter Trust

Special Thanks:
- Data provided by Wiebke Koepp, Tanny Andrea Chavez Esparza, Alexander Hexemer, and Dula Parkinson, Advanced Light Source, LBNL
- Brendt Wohlberg, LANL;
- Venkat (Singanallur) Venkatakrishnan, ORNL;
- Xiao Wang, ORNL