Image Processing with Nonparametric Neighborhood Statistics

Ross T. Whitaker

Scientific Computing and Imaging Institute
School of Computing
University of Utah

Suyash Awate, University of Pennsylvania, Dept of Radiology

See also: http://www.cs.utah.edu/~suyash/pubs/uinta.html

Talk Overview

- Motivation
- Image denoising
- Density estimation
- UINTA filtering strategy overview
- Entropy minimization
- Implementation issues: statistics, image processing
- Other microscopy work
- Final thoughts

Images

Denoising Vs Reconstruction

- Any geometric/statistical penalty can be applied in two ways:
  1. Gradient descent as filter (choose # iterations)
  2. With data (fidelity) term to steady state
     - Variational
     - Noise/measurement models, optimality, etc.

Variational Methods
E.g Anisotropic Diffusion

- Perona&Malik (1990)

$$\frac{\partial f}{\partial t} = \nabla \cdot c(|\nabla f|)\nabla f$$

- Penalty:
  - Quadratic on grad-mag with outliers (discontinuities)
    - Nordstrom 1990; Black et. al 1998
  - Favors piecewise const. Images

Other Flattening Approaches

- Total variation
  - Rudin et. al (1992)
- Mumford-Shah (1989) related
  - Explicit model of edges
  - Cartoon model
- Level sets to model edges
  - Chan & Vese (2000)
  - Tsai, Yezzi, Willsky (2000)
- Model textures + boundaries
  - Meyer (2000)
  - Vese & Osher (2002)

PDE Methods
Other Examples

- Weickert (1998)
  - Coherence enhancing

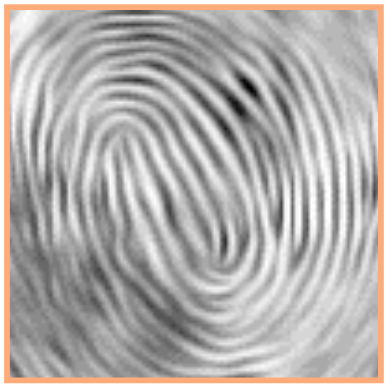

- Tasdizen et. al (2001)
  - Piecewise-flat normals

- Wilmore flows
  - Minimize curvature

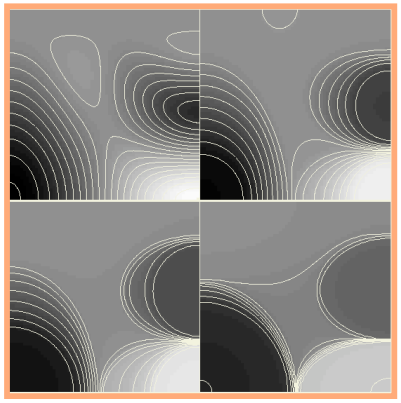

Issues

- Prioritize geometric configurations a priori
  - Works well of the model fits, otherwise...
- Free parameters
  - Thresholds -> determine when to apply different models (e.g. "preserve edge or smooth")
- Generality
  - Cartoon-like simplifications are disastrous in many applications
- Increasing the geometric complexity
  - Is there a better way?

Examples

MRI (Simulated noise)

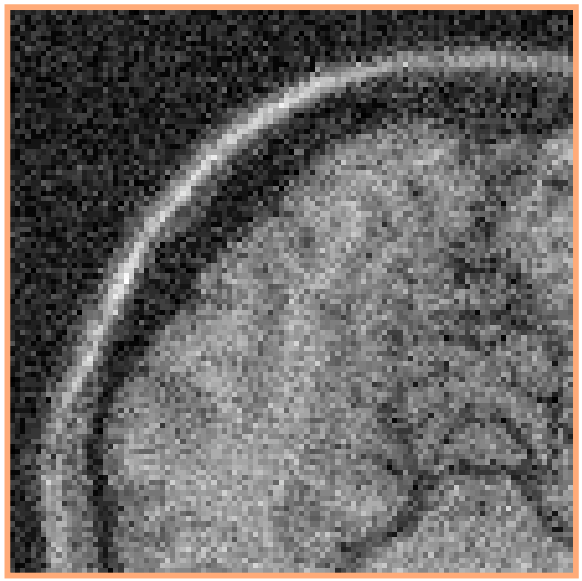

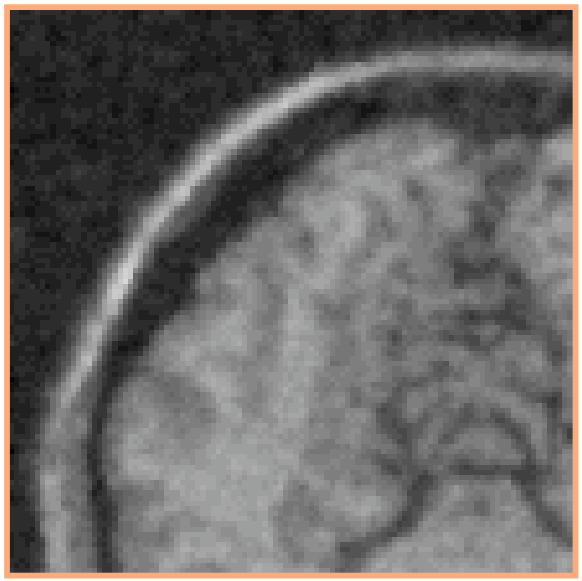

Bilateral Filtering

Coherence Enhancing

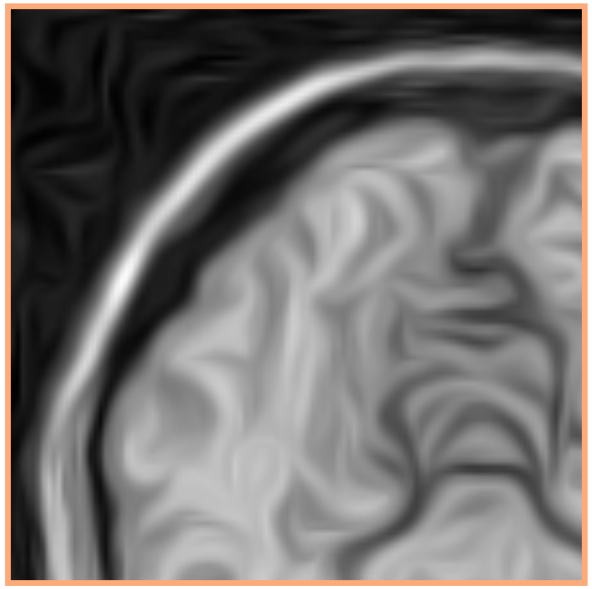

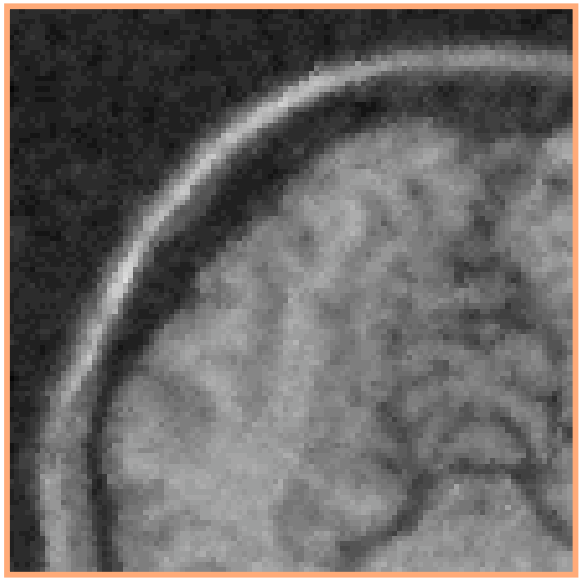

Anisotropic Diffusion

Observations About Images

- Statistics of *natural* images are not *so* random
  - Huang & Mumford (1999)
- But not so simple
  - Manifolds in high-dimensional spaces
  - de Silva & Carlsson (2003)

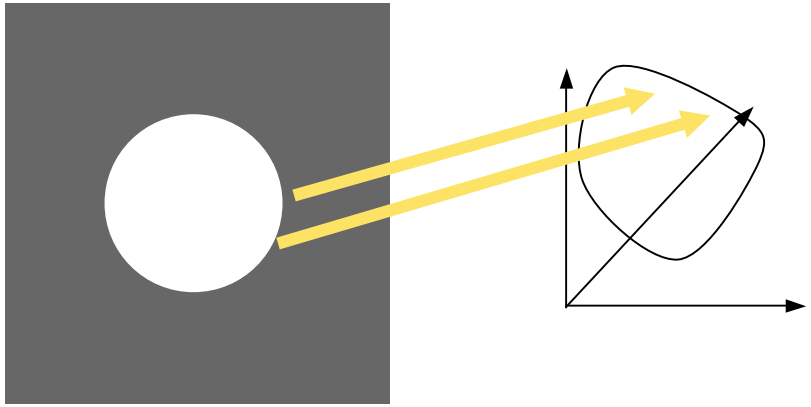

Related Work

- DUDE algortihm-Weissman et. al (2003)
  - Discrete channels + noise model
  - MLE estimation
- Texture synthesis
  - Efros & Leung (1999)
  - Wei & Levoy (2002)
- NL-means, Baudes et al. (CVPR 2005)
  - Independent, simultaneously presented
  - More later…
- Sparsity in image neighborhoods
  - Roth and Black 2005
  - Elad and Aharon 2006

Image Model

- Pixels and neighborhoods Z = (X, Y)
  - P(Z), P(X|Y)
- Scenario
  - Corrupted image -> noise model
  - Prior knowledge P(X|Y)
  - Theorems:
    - Can produce most likely image x' using P(X|Y = y')
    - Iterate to produce optimal estimate

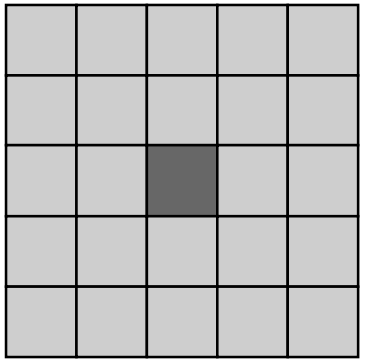

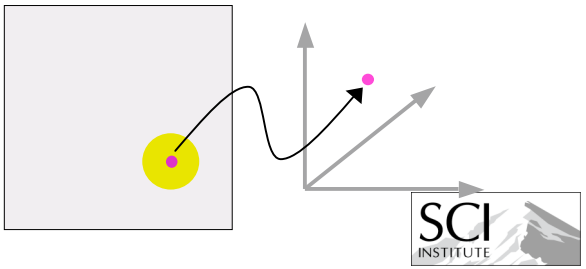

Modeling P(Z)

- Set of image neighborhoods
  - Large, complex, high-dimensions
- Approach
  - Represent complexity through examples
  - Nonparametric density estimation

Nonparametric, Multivariate Density Estimation

- **Nonparametric estimation**
  - No prior knowledge of densities
  - Can model *real* densities
- **Statistics in higher dimensions**
  - Curse of dimensionality (volume of n-sphere -> 0)
  - + However, empirically more optimistic
  - + Z has identical marginal distributions
  - + Lower dimensional manifolds in feature space

Parzen Windows (Parzen 1962)

- Scattered-data interpolation

$$p(z) \approx \frac{1}{|A|} \sum_{z_i \in A} G(z - z_i, \psi)$$

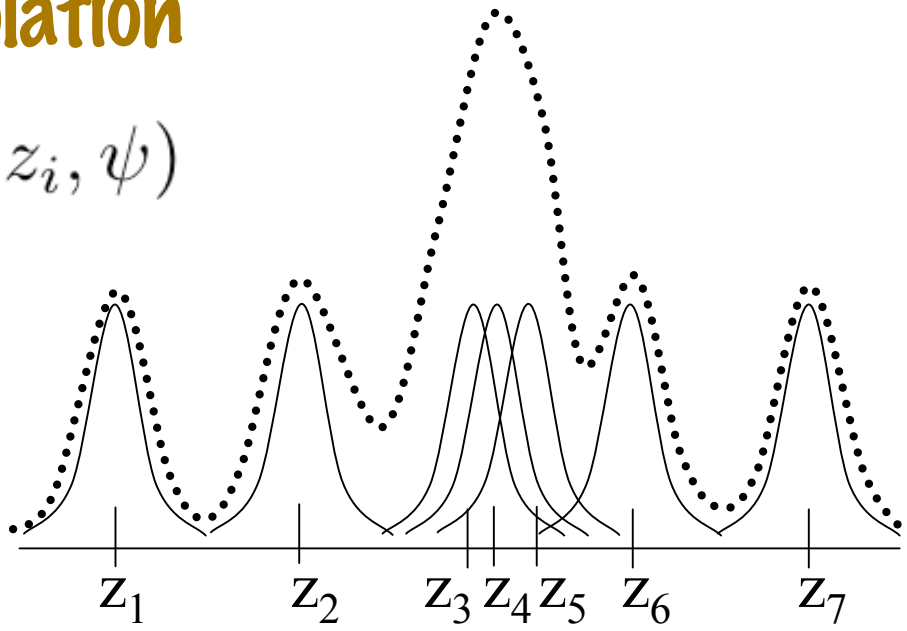

- Window function
  - $G \equiv$ Gaussian
  - Covariance matrix: $\psi = \sigma^2 I$

Entropy (Shannon 1948)

- Entropy of a random variable *X* (instance x)
  - Measure of *uncertainty* – information content of a sample

$$h(X) = -\int p(x) \log p(x) dx = -E_p \left[\log p(X)\right]$$

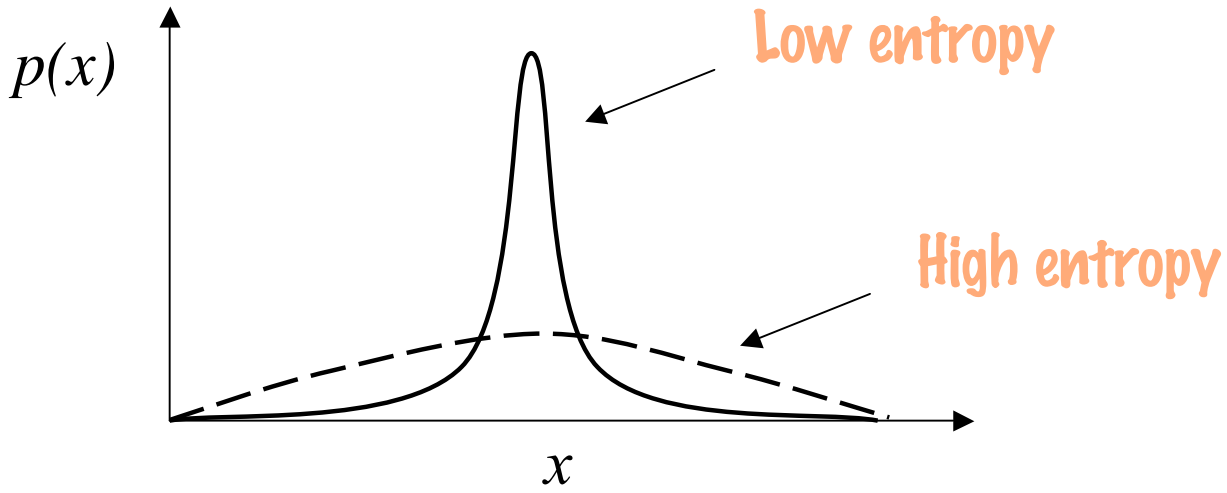

UINTA Strategy

Awate & Whitaker CVPR 2005, PAMI 2006

- Iterative algorithm
- Progressively minimizes the entropy of image nhds Z = (X, Y)
  - Pixel entropies (X) conditioned on nhd values (Y)
  - Gradient descent (time steps -> mean shift)
- Nonparametric density estimation
  - Stochastic gradient descent

Entropy Minimization

- Entropy as sample mean

$$\begin{aligned} h(Z) &= -E_p[\log p(Z)] \\ &\approx \frac{1}{|B|}\sum_{i\in B}\log p\left(z_i\right) \\ &\approx \frac{1}{|B|}\sum_{i\in B}\log\left(\frac{1}{|A|}\sum_{j\in A}G(z_i - z_j, \psi)\right) \end{aligned}$$

  - Set *B*: all pixels in image
  - Set *A*: a small *random* selection of pixels
  - z_i shorthand for $z(s_i)$
- Stochastic approximation

Entropy Minimization

- Stochastic approximation
  - Reduce $O(|B|^2)$ to $O(|A||B|)$
  - Efficient optimization
- Stochastic-gradient descent

$$\begin{aligned}\Delta x &= -\lambda \frac{\partial h(X|Y=y)}{\partial x} \\ &\approx \frac{\lambda \psi^{-1}}{|B|}\left[\sum_{j\in A} \frac{G(z_j - z, \Psi)}{\sum_{k\in A} G(z_k - z, \Psi)} x_j - x\right]\end{aligned}$$

SCI INSTITUTE

Mean-Shift Procedure (Fukunaga et al. 1975)

- Entropy minization <-> mean shift

$$\lambda = \Psi|B| \quad x \longleftarrow \sum_j w_j x_j$$

- Mean-shift - a mode seeking procedure

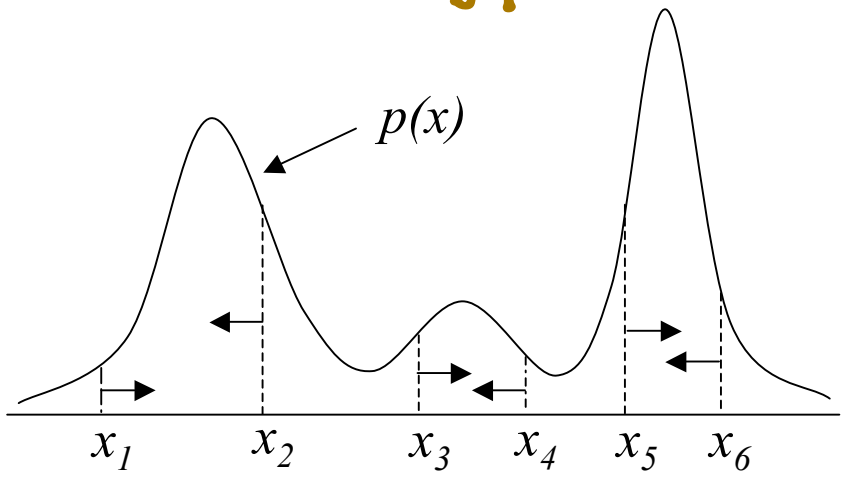

SCI INSTITUTE

Mean-Shift Procedure (Fukunaga et al. 1975)

- Data filtering to reduce noise
  - Hand tuned parameters

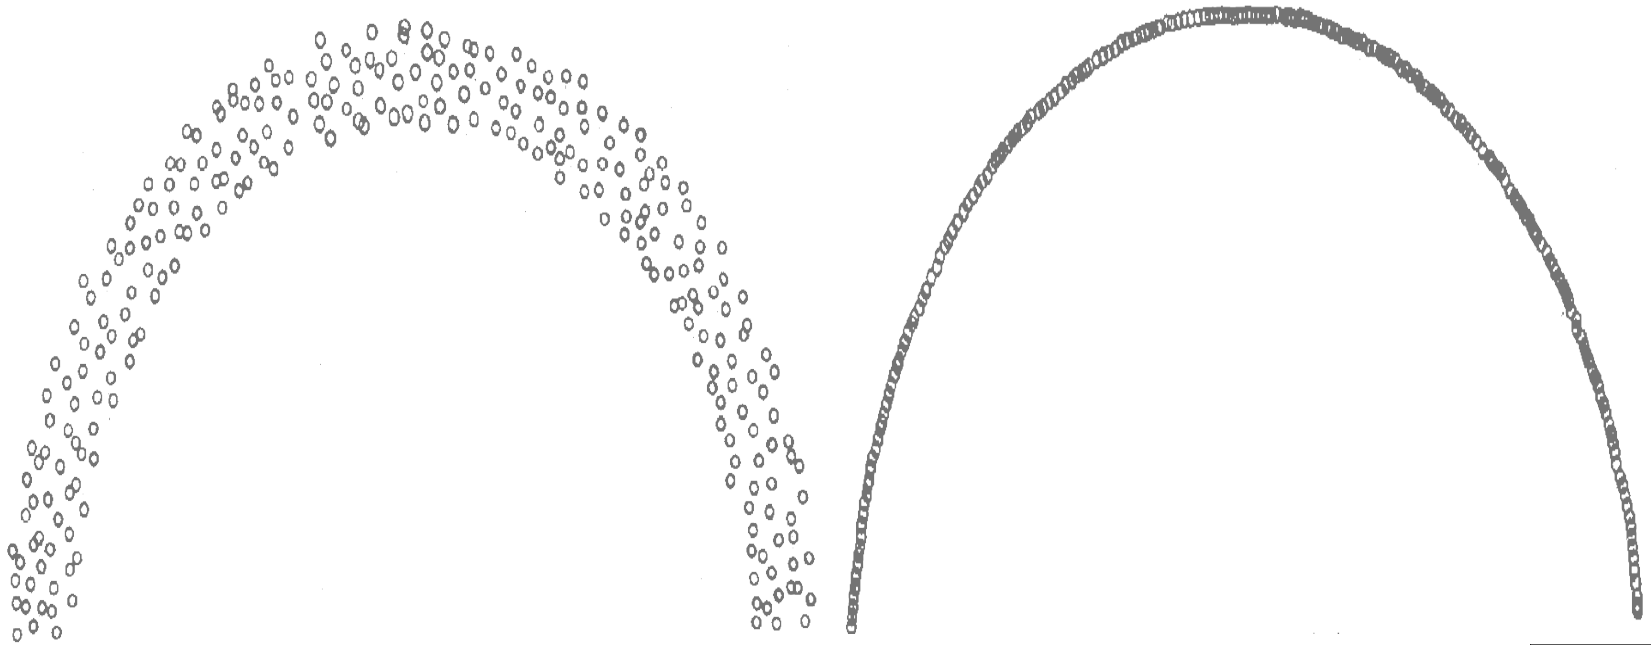

SCI INSTITUTE

Implementation Issues

- Scale selection for Parzen windowing
  - Automatic – min entropy with cross validation
- Rotational invariance
- Boundary neighborhoods
- Random sample selection – nonstationary image statistics
- Stopping criteria

Results

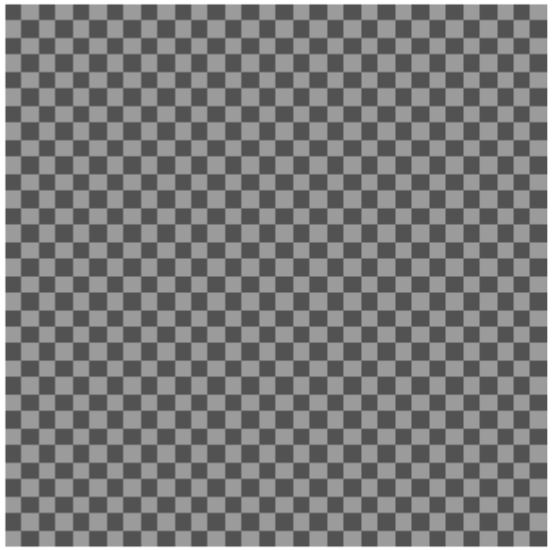

Original

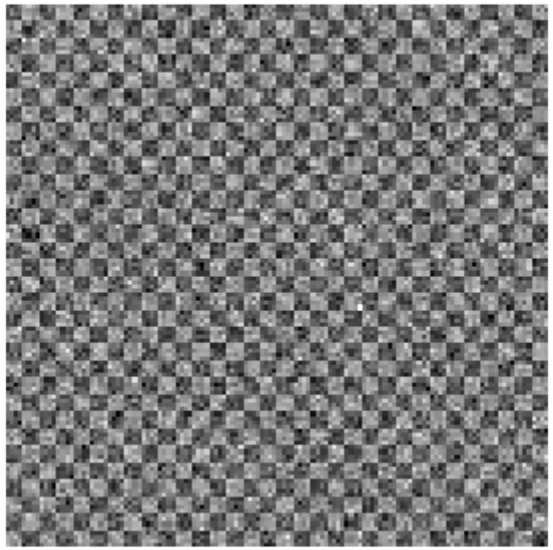

Noisy

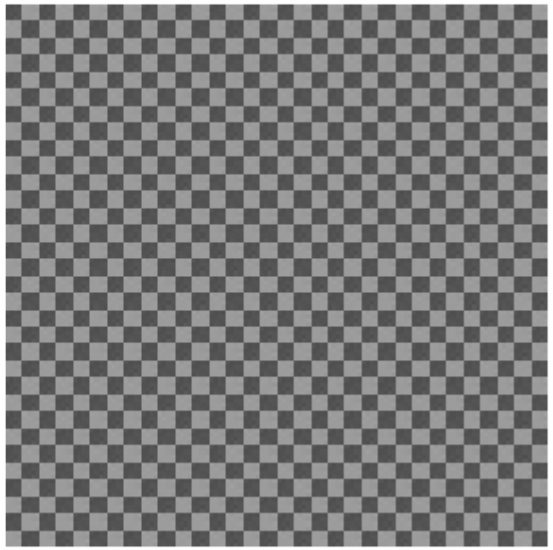

Filtered

Checkerboard With Noise

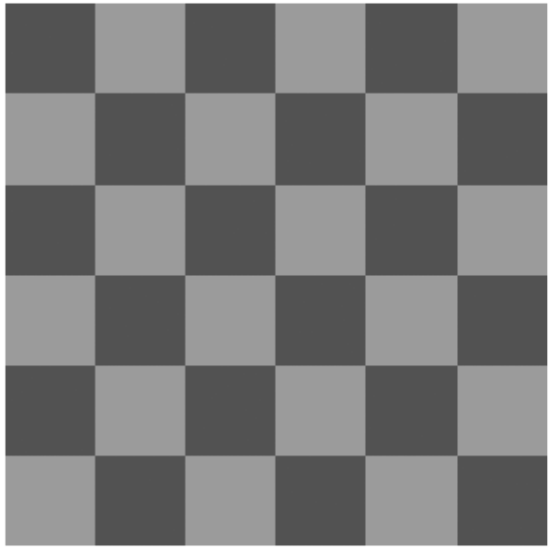

Original

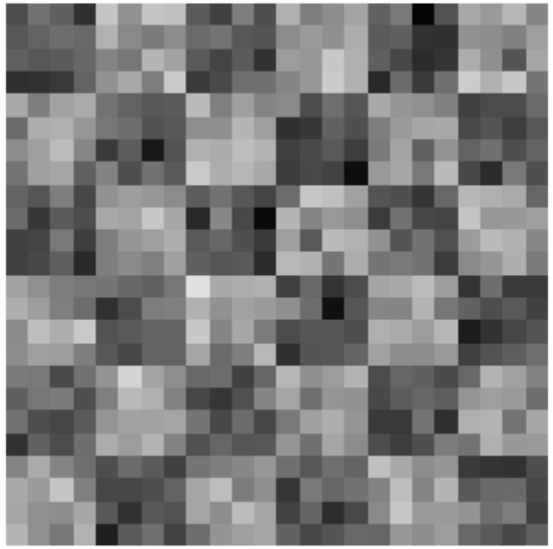

Noisy

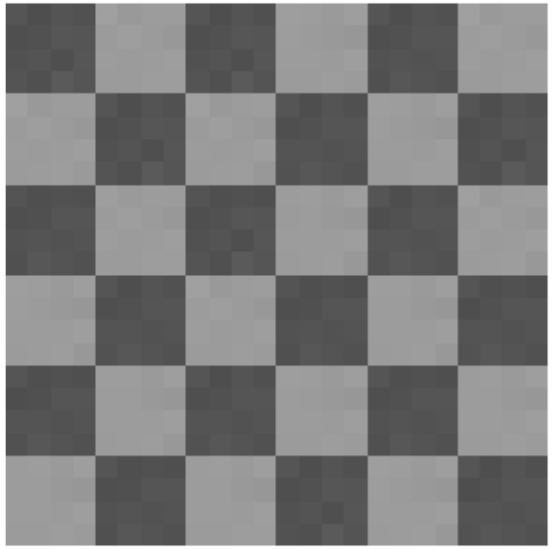

Filtered

Quality of Denoising

- σ, joint entropy, and RMS- error vs. number of iterations

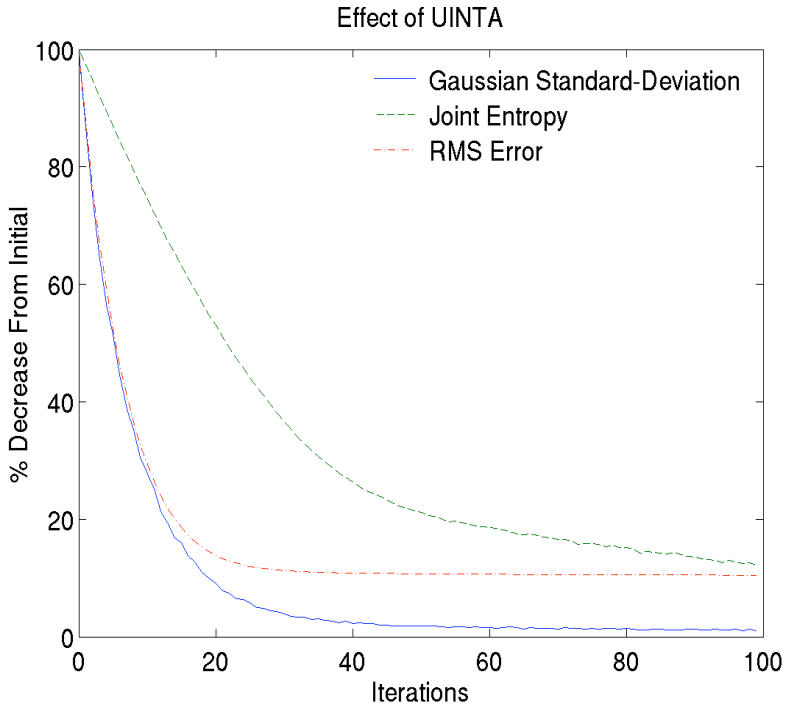

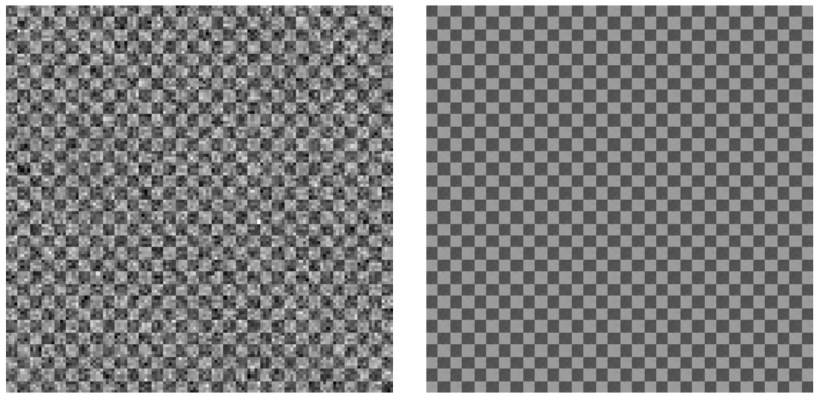

Vs Perona Malik

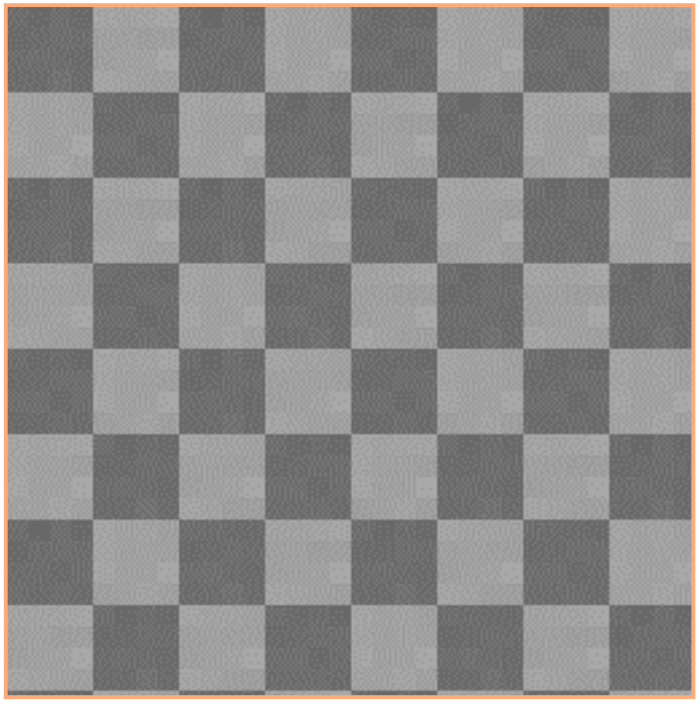

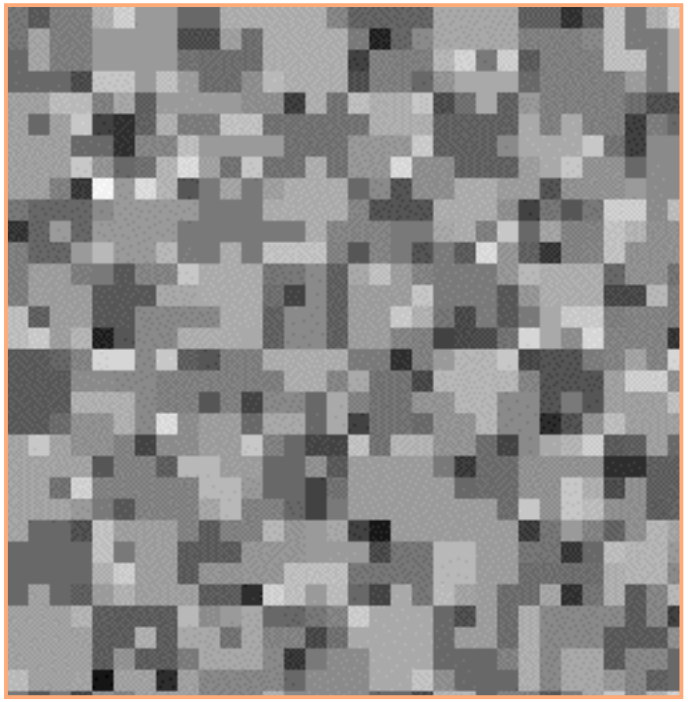

MRI Head

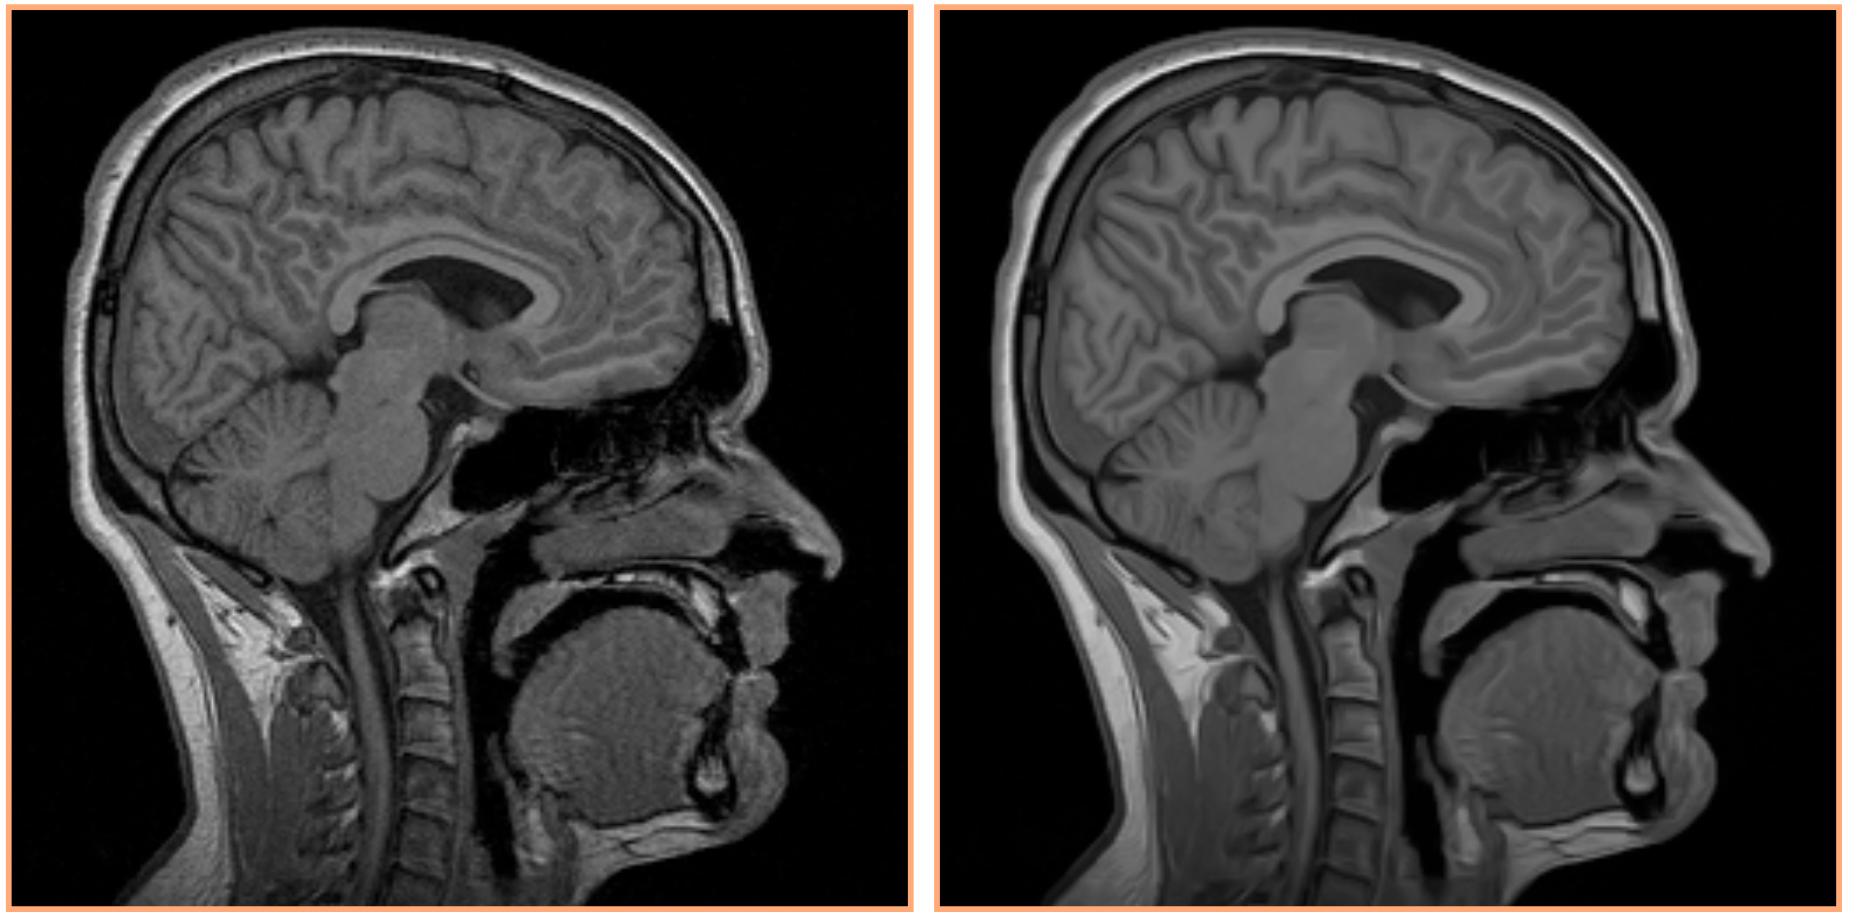

MRI Head

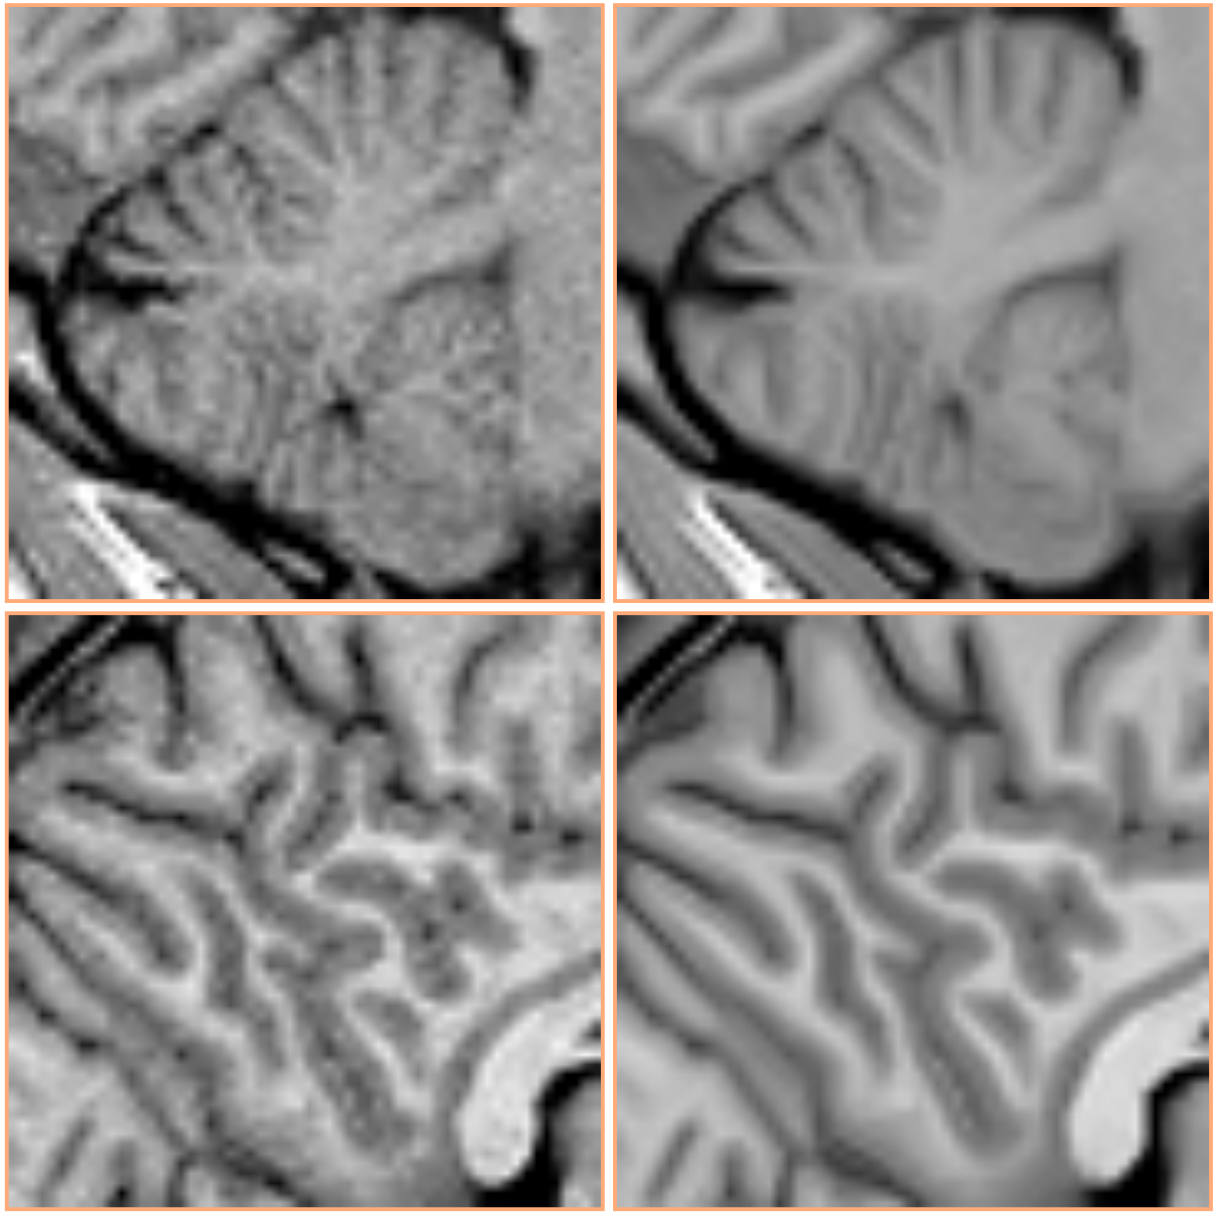

SCI INSTITUTE

Fingerprint

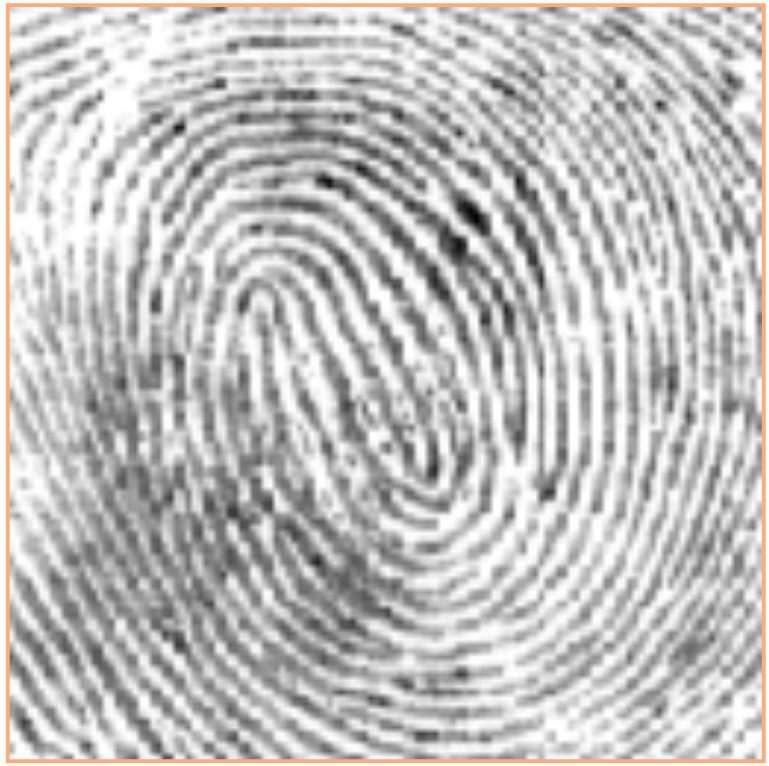

Fingerprint

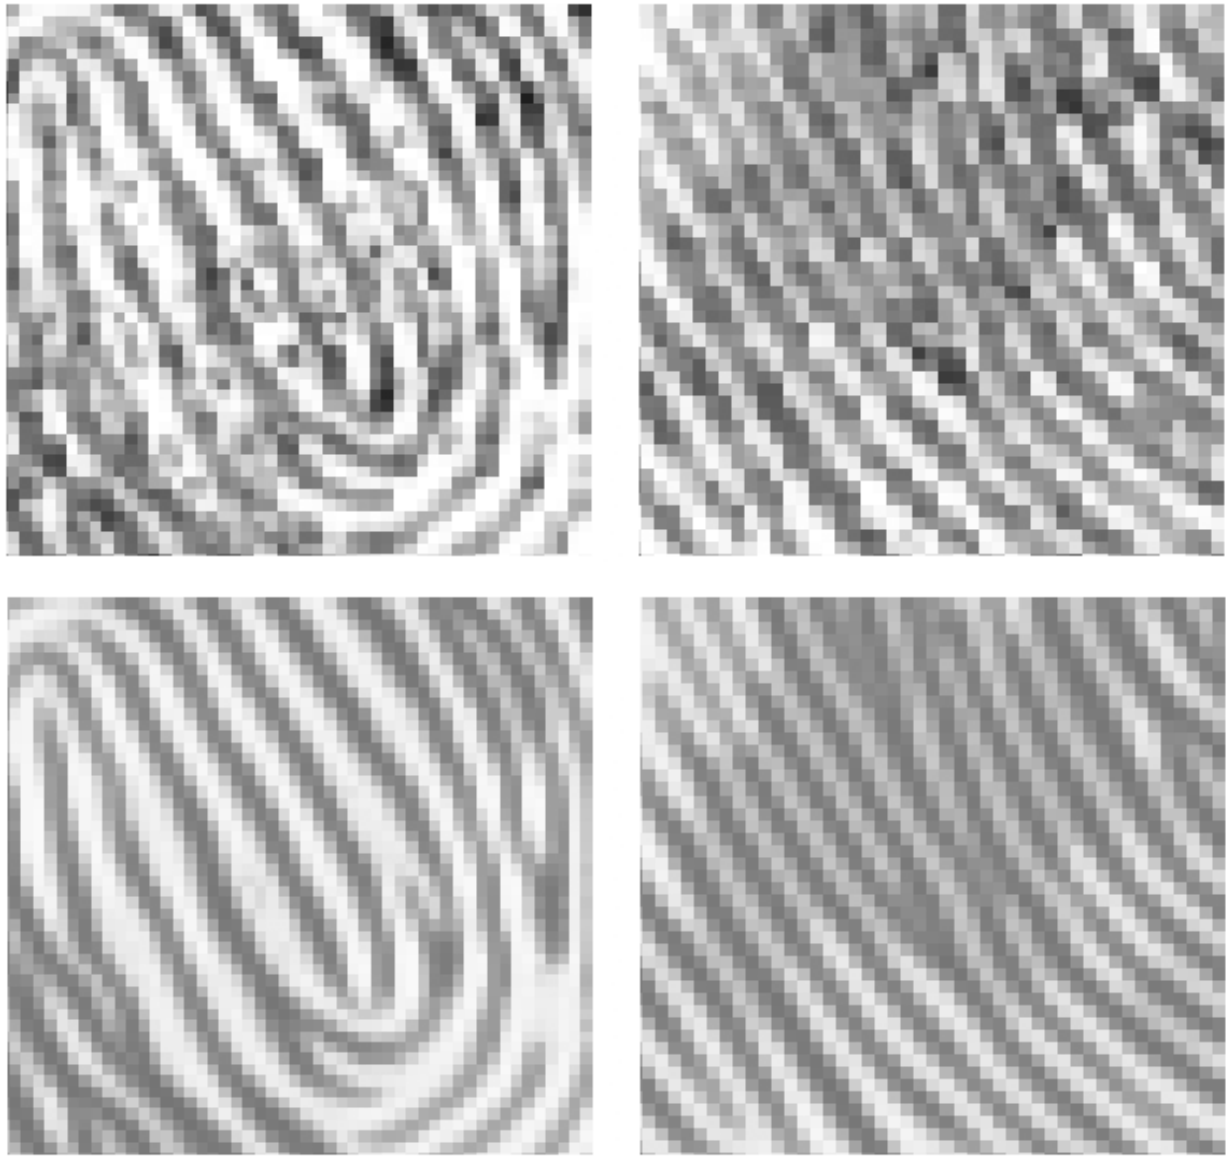

SCI INSTITUTE

Vs Perona Malik

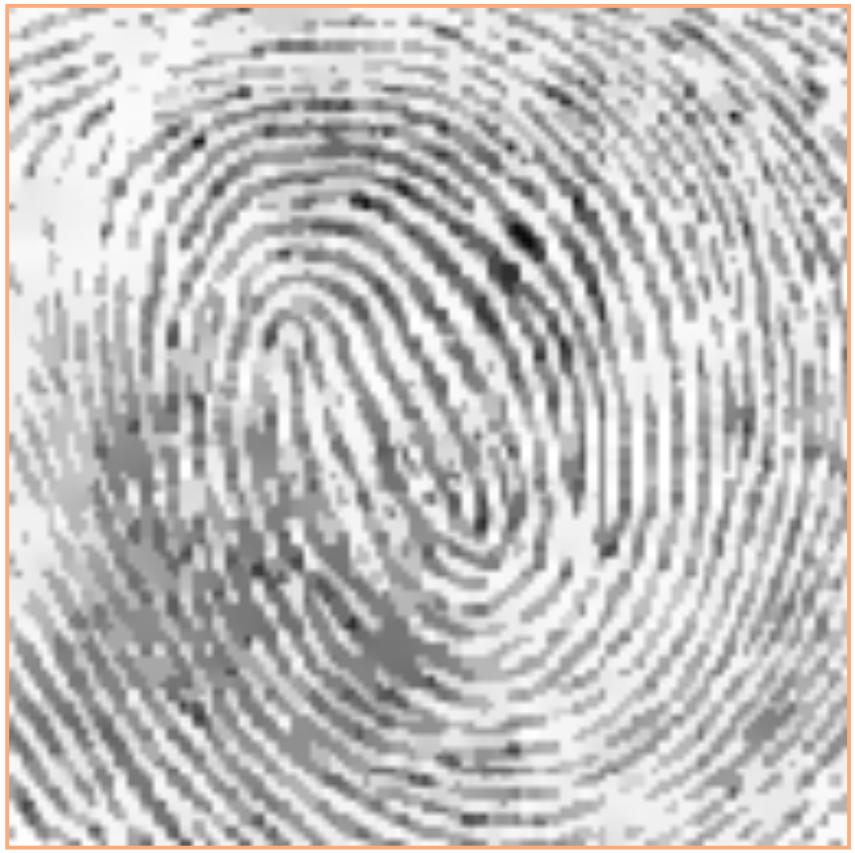

Vs Coherence Enhancing

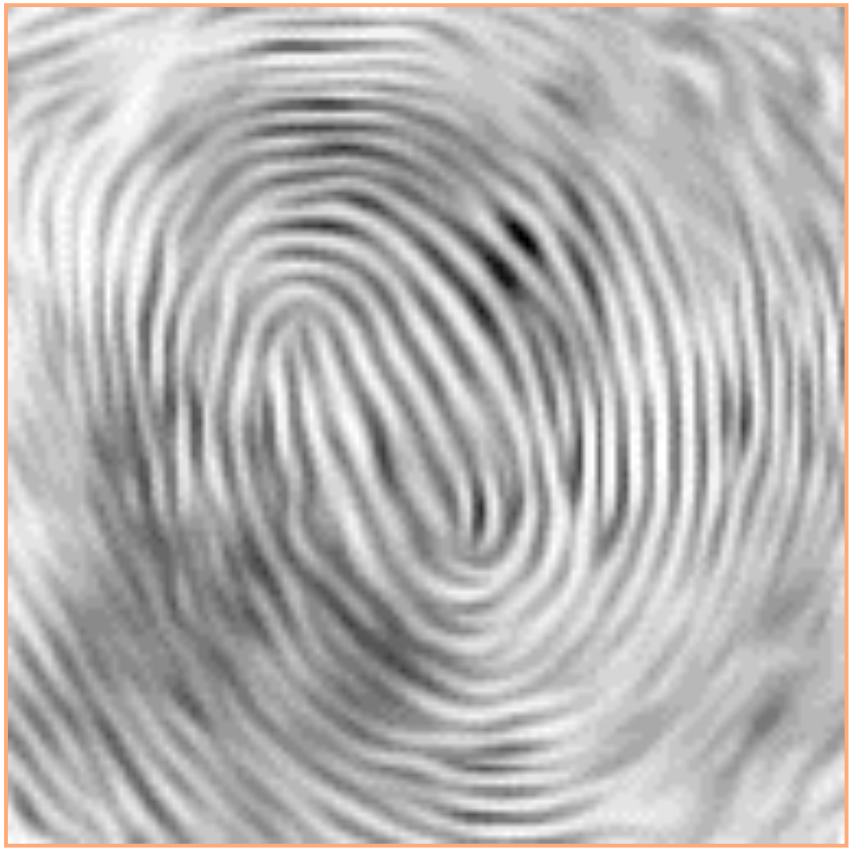

Lena

SCI INSTITUTE

Lena

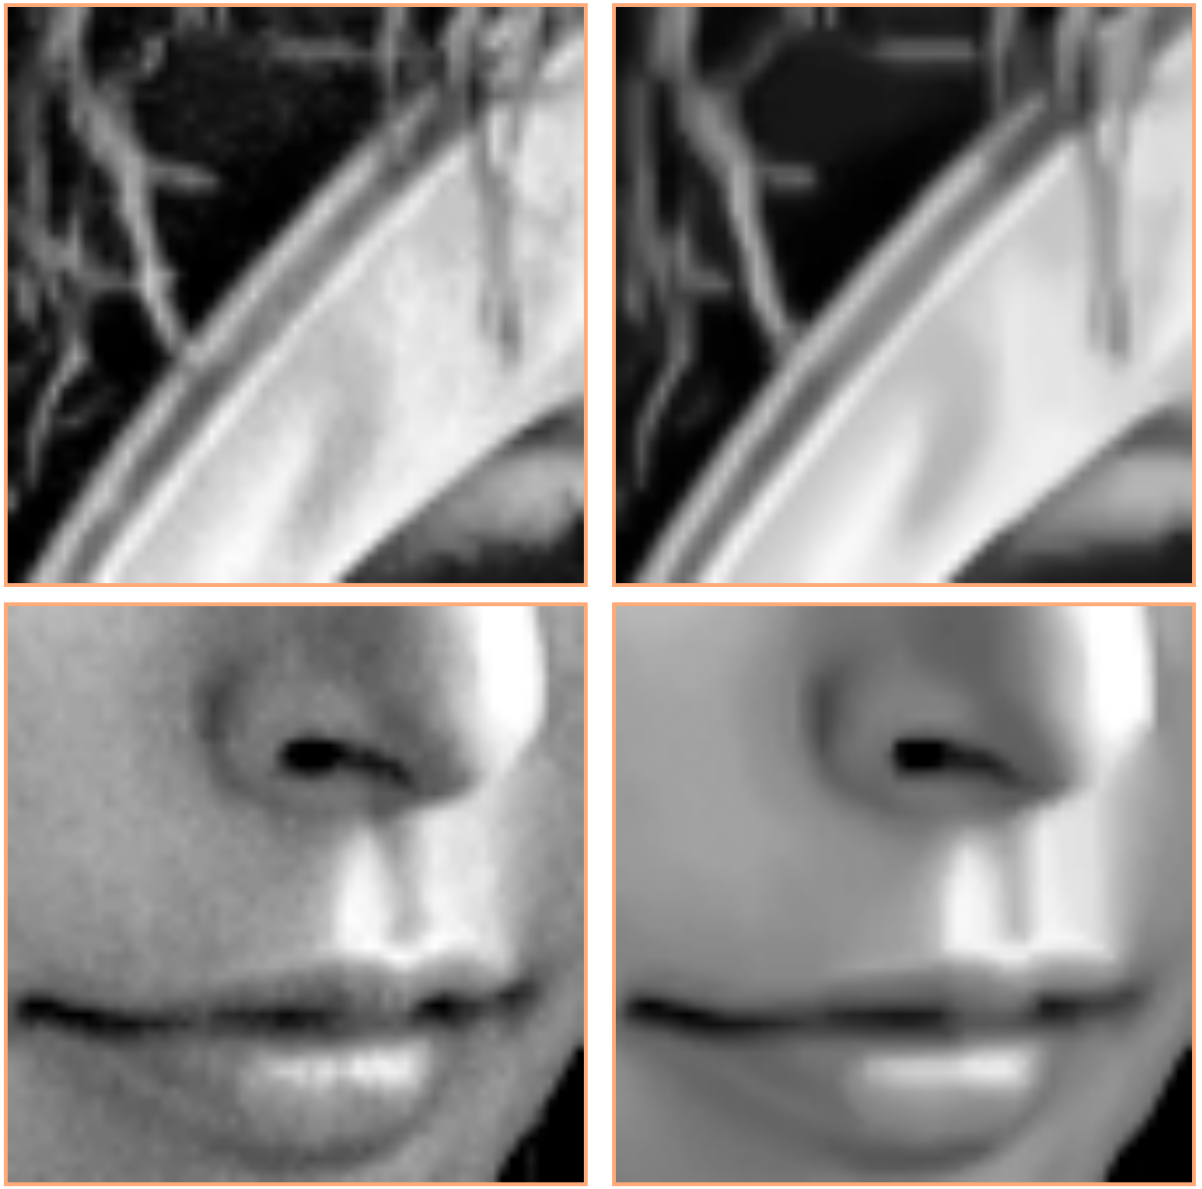

Results

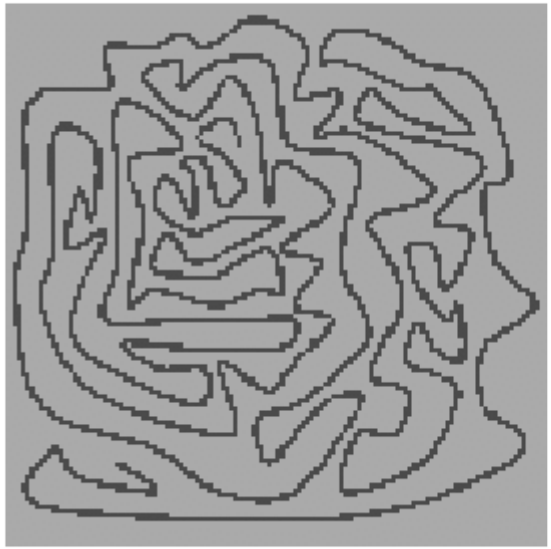

Original

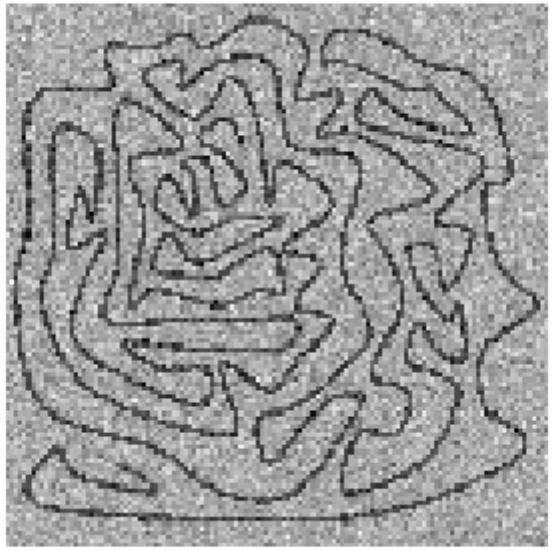

Noisy

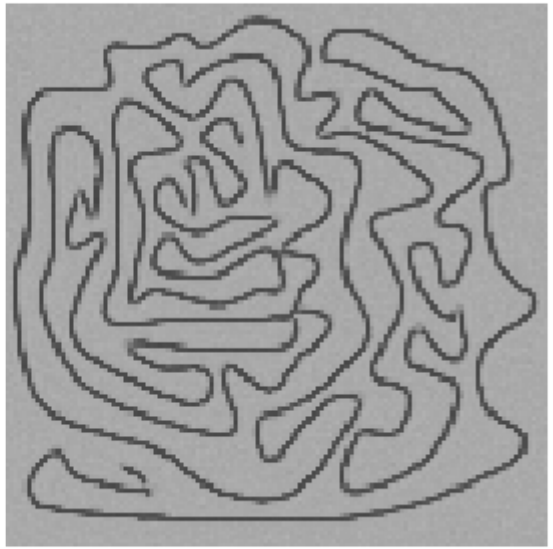

Filtered

Results

Original

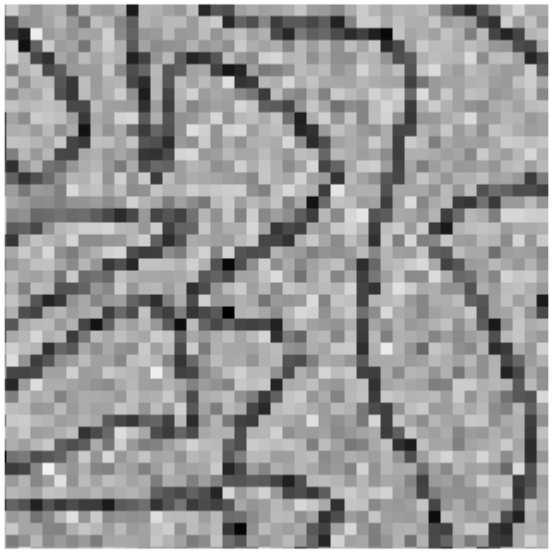

Noisy

Filtered

Results

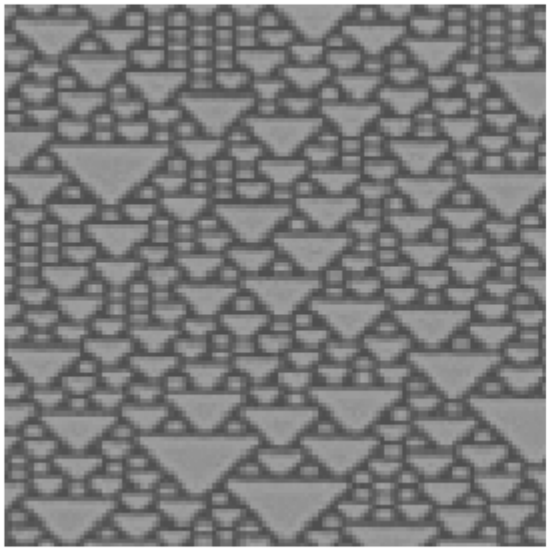

Original

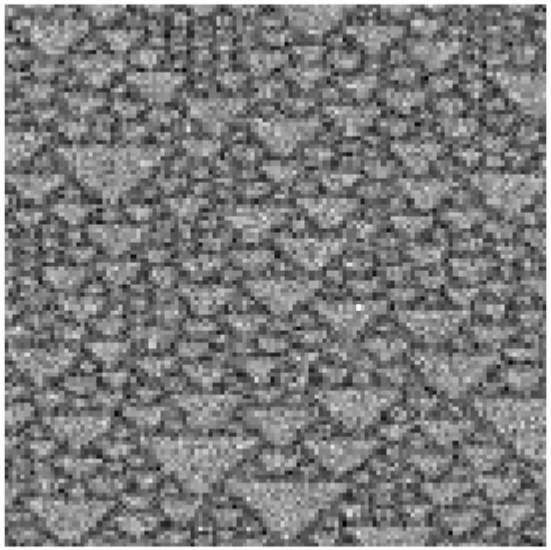

Noisy

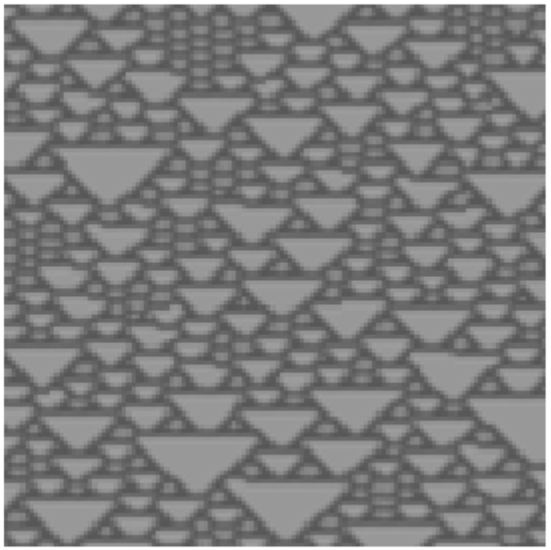

Filtered

Fractal

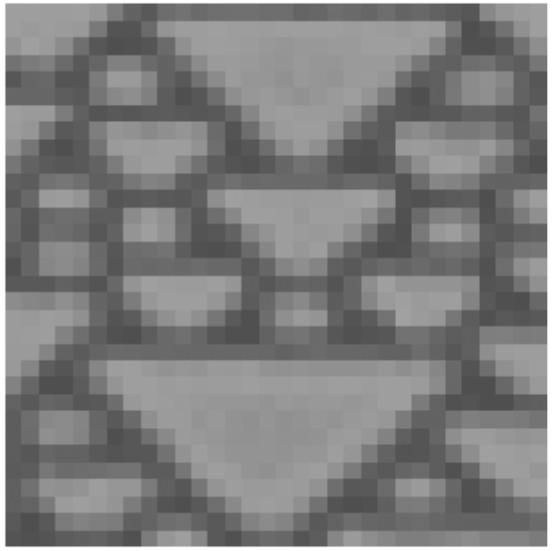

Original

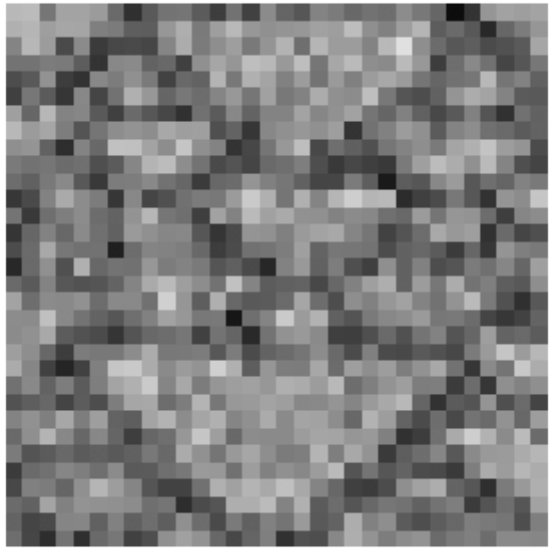

Noisy

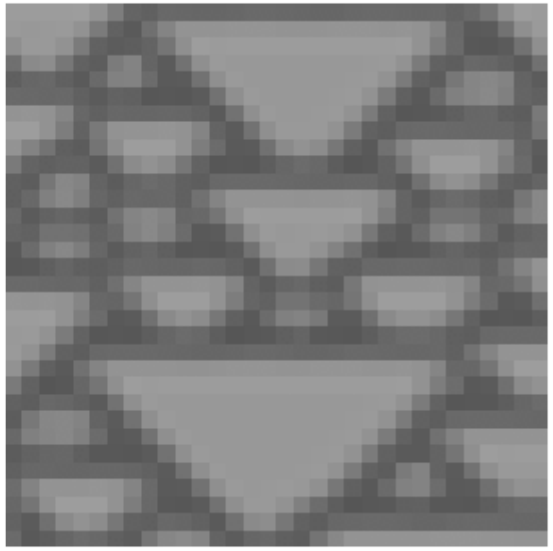

Filtered

Microscopy

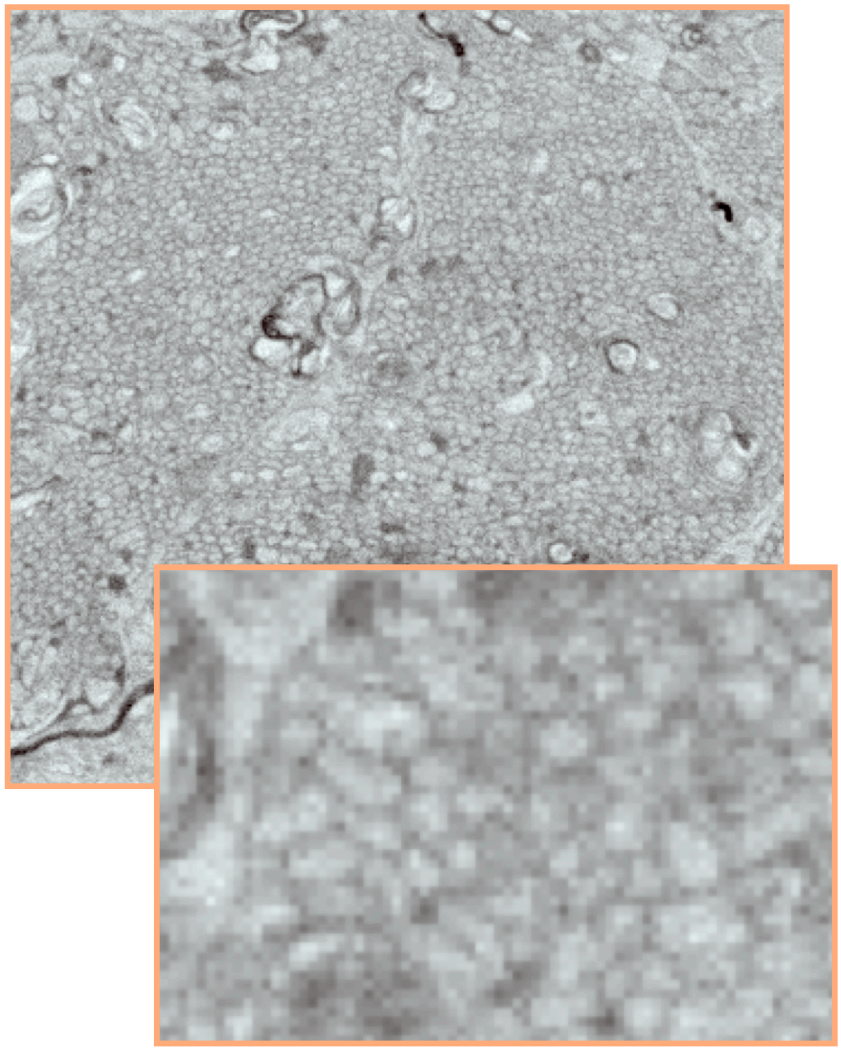

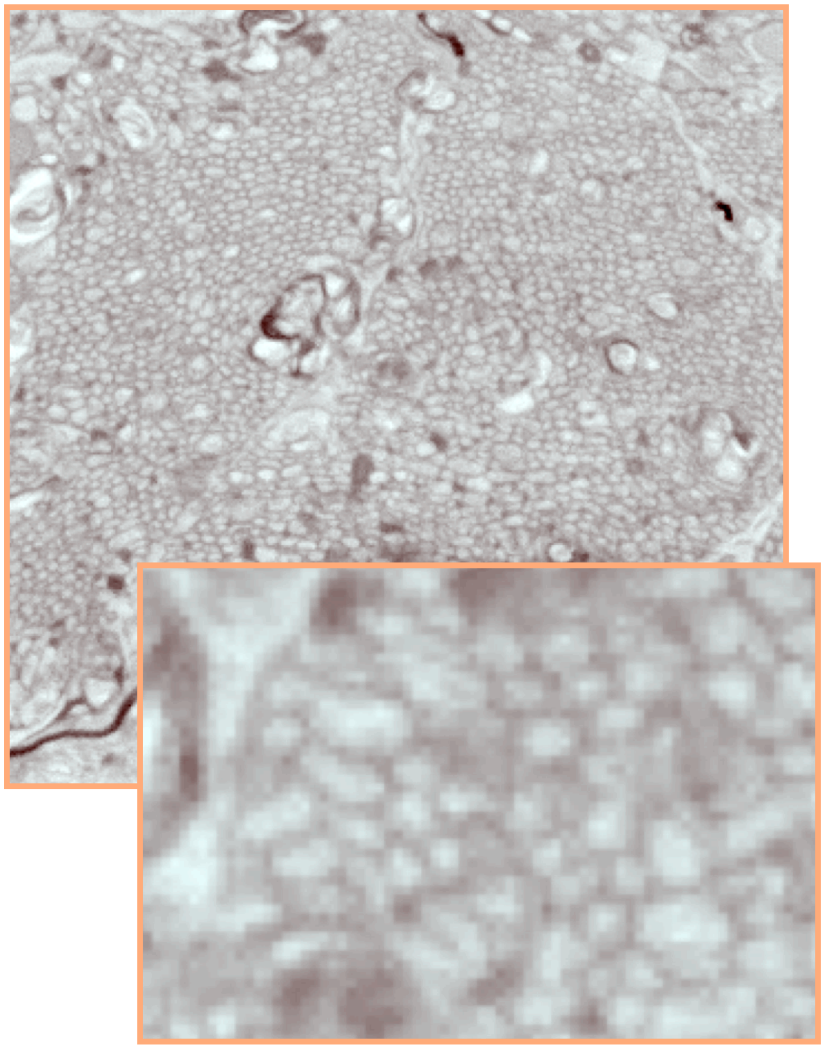

Quantitative Results

- Generalizes well
  - Relatively insensitive to a few parameters (e.g. nhd size)
- Compares favorably with s.o.t.a. wavelet denoisers
  - Close but worse for standard images (photographs)
  - Better for less typical images (defy wavelet shrinkage assumptions)
- Spectral data -> gets even better

Other Applications

- Optimal estimation/reconstruction
  - IPMI 05, TMI 07

noiseless | Rician noise | *estimated* prior | reconstructed

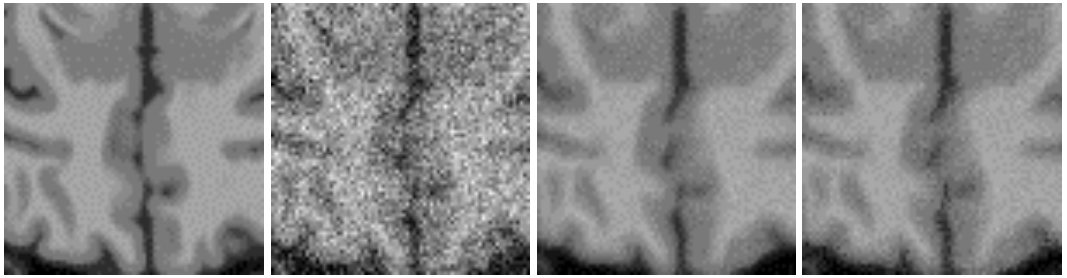

- Tissue classification
  - MICCAI 05, MedIA 06

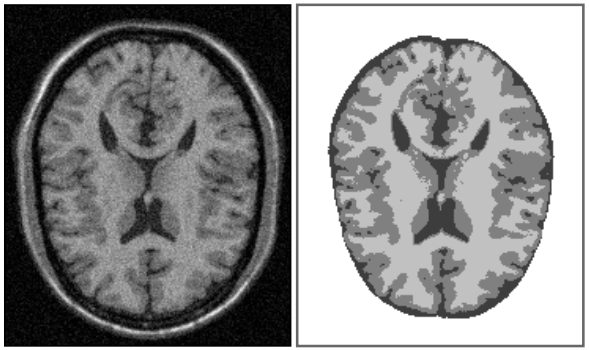

- Segmentation
  - ECCV 05

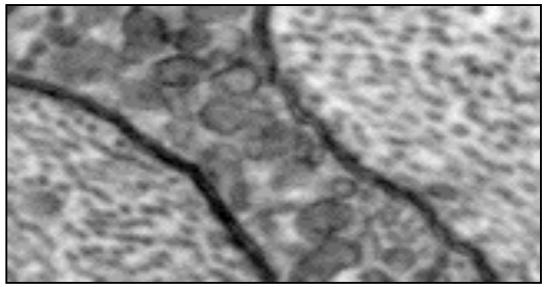

Other Work in Microscopy

ET Surface Reconstruction

- Limited-angle tomography artifacts
  - Varies with recon technique
- Approximate solution
  - Smooth with discontinuity at interface
  - E.g. anatomical boundary
- Fit model directly to tilt-series data
  - Refine interface iteratively
  - Deformable model

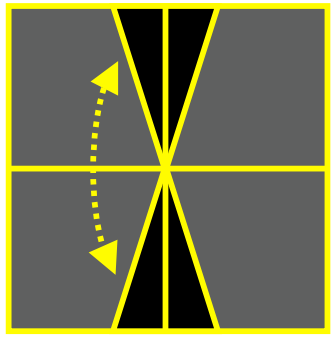

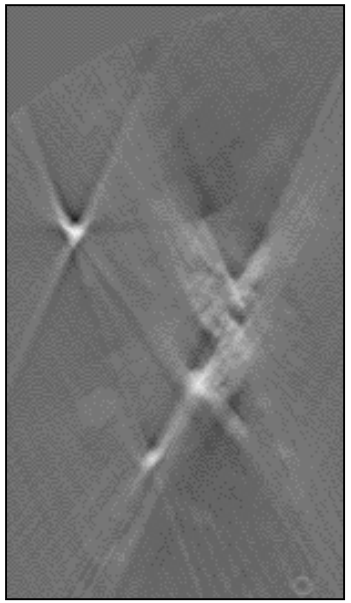

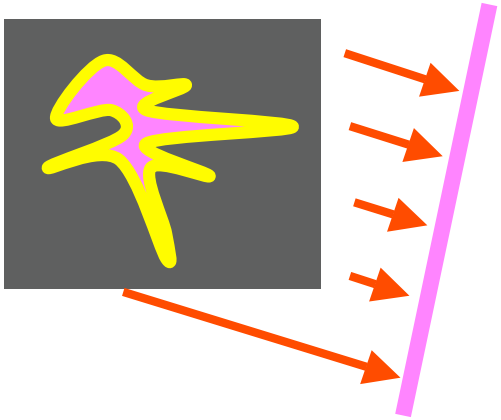

ET Surface Reconstruction

Initialization (BP)

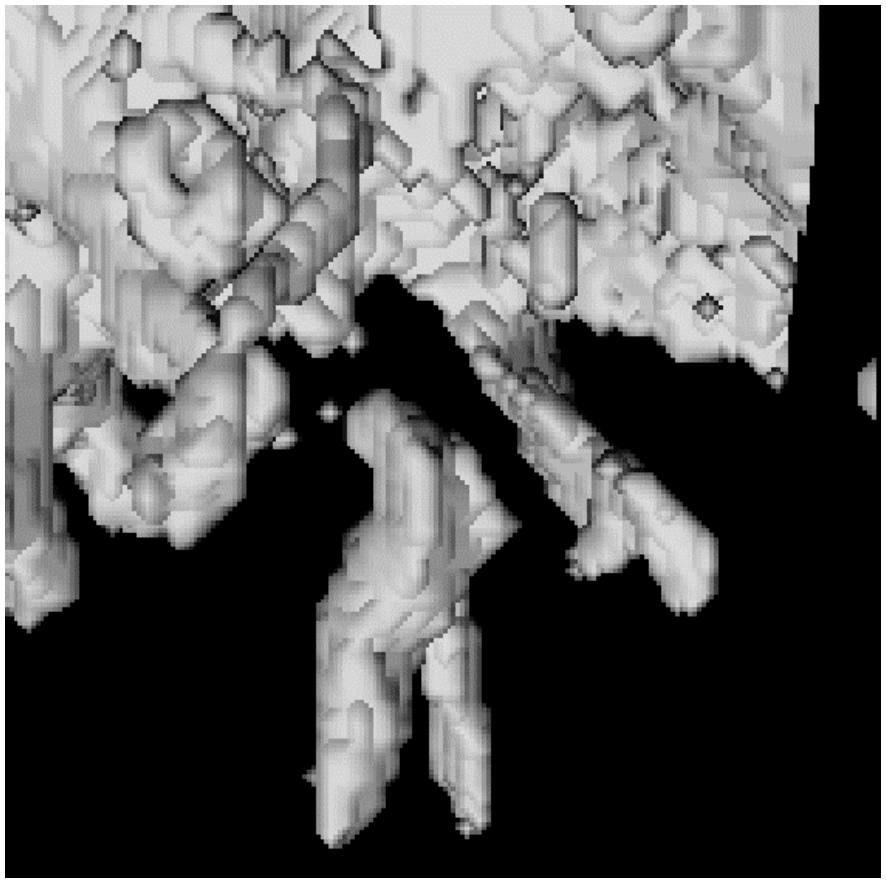

Final Reconstruction

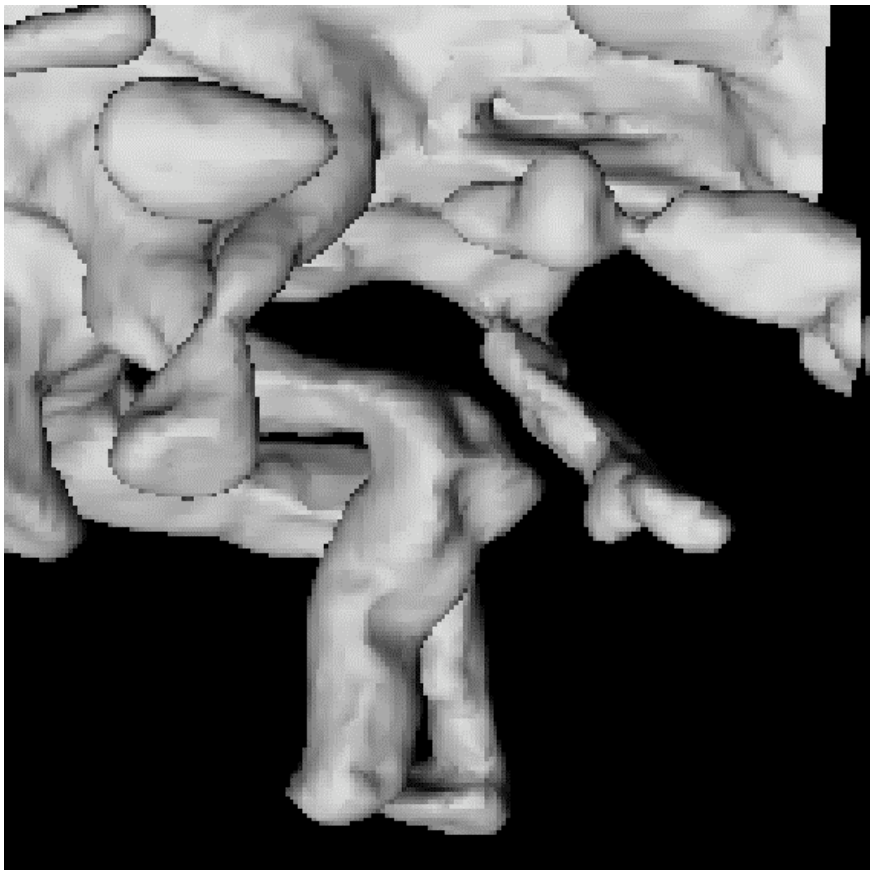

Elangovan & Whitaker 2001, Whitaker & Elangovan 2002

SCI INSTITUTE

Interactive 3D Tools

- Volume rendering

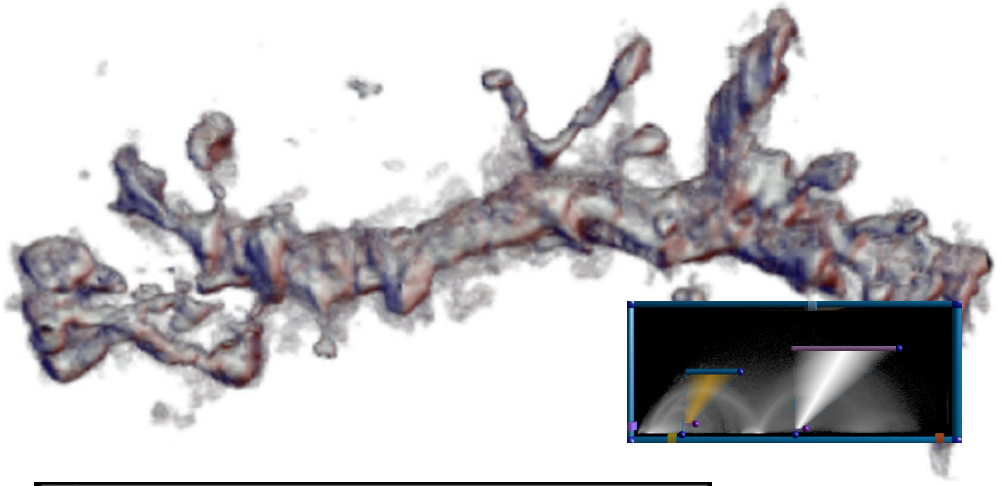

- Seg3D
  - www.seg3d.org

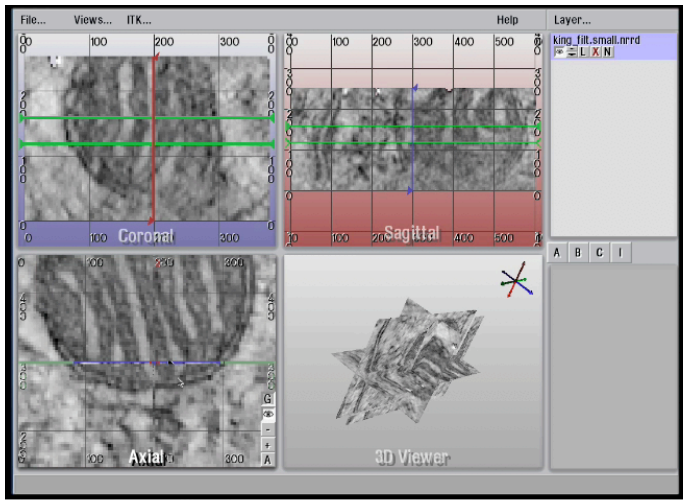

SCI INSTITUTE

Retinal Mapping -Marc, Tasdizen

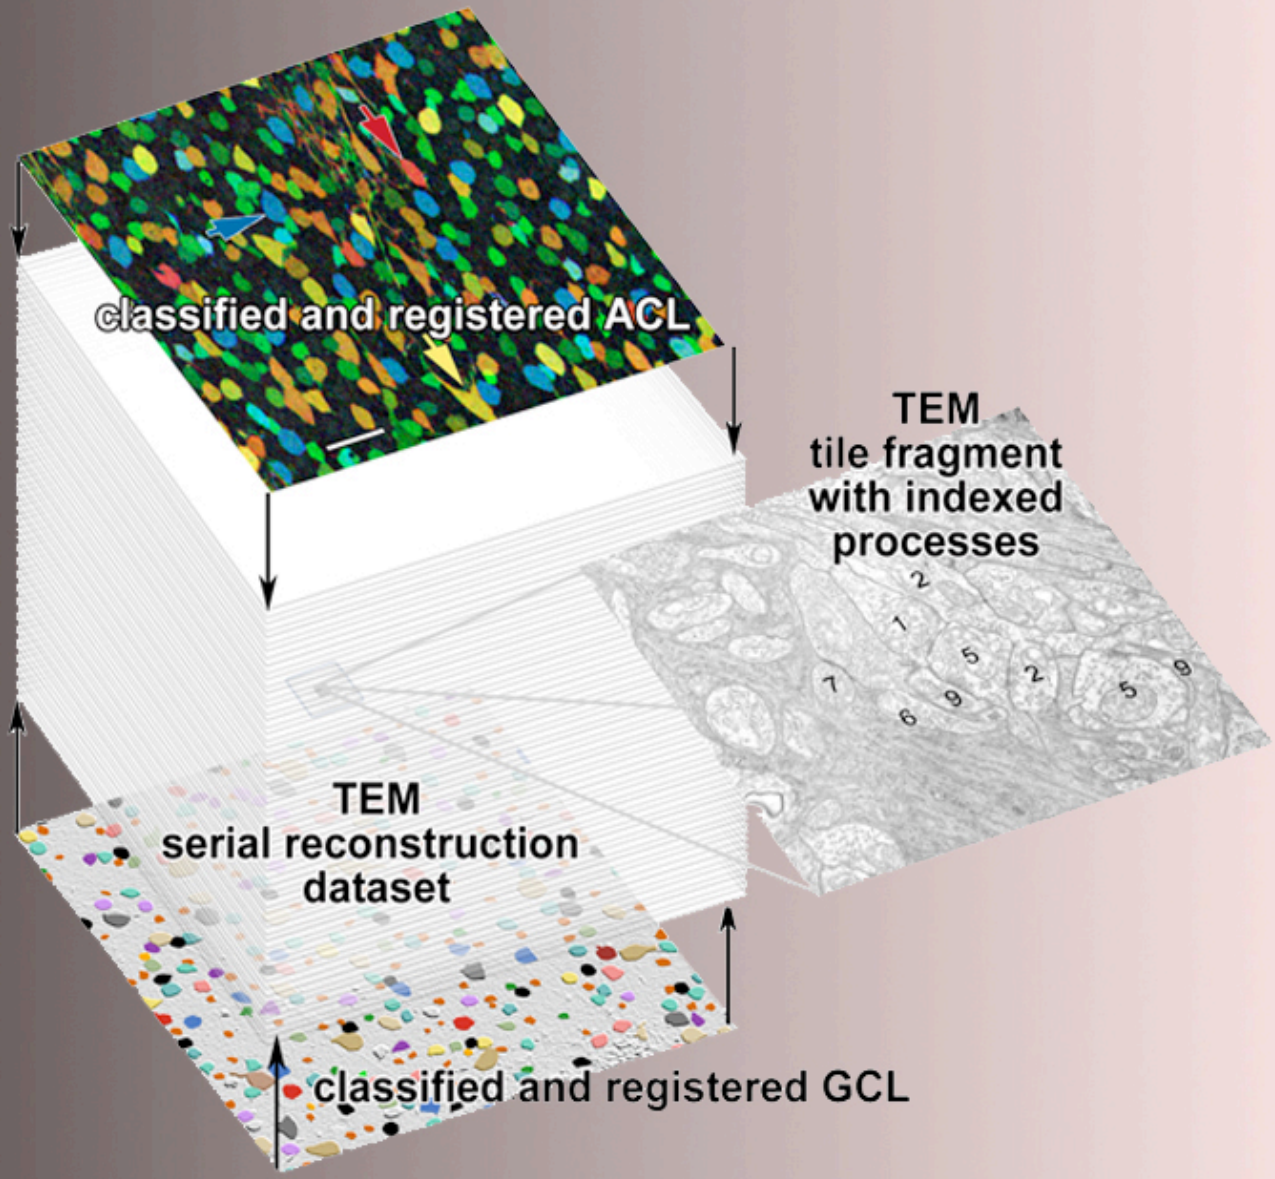

Tiling Challenges

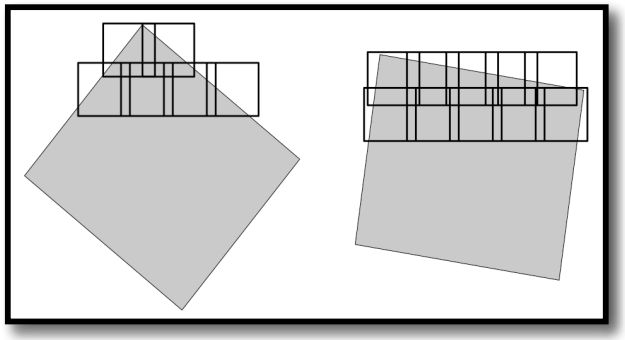

Tile arrangement

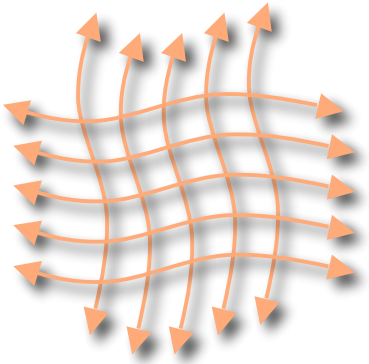

Warping

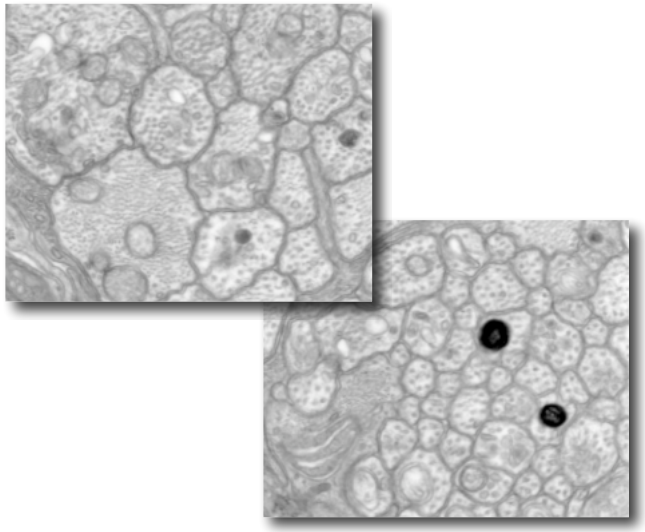

Data overlap/complexity

Final Stack

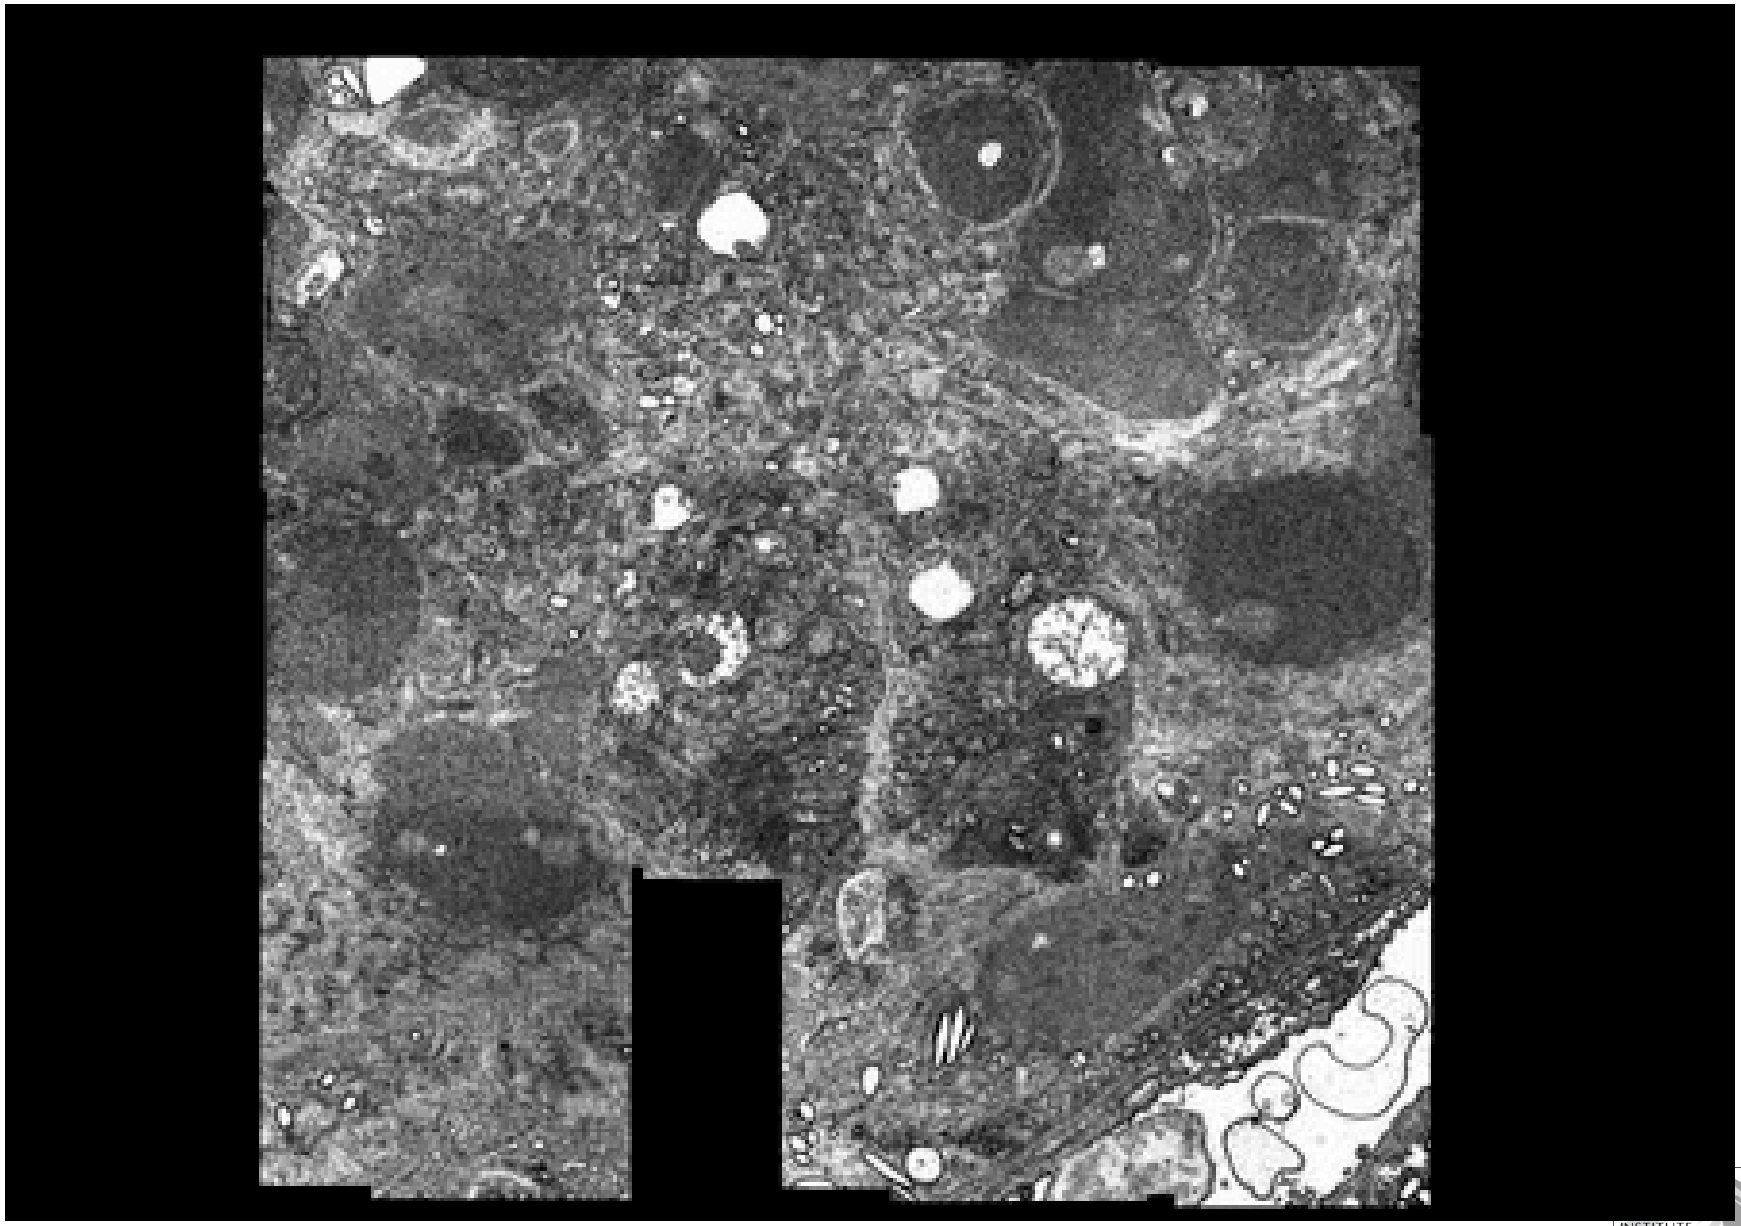

C. Elegans - Jorgensen

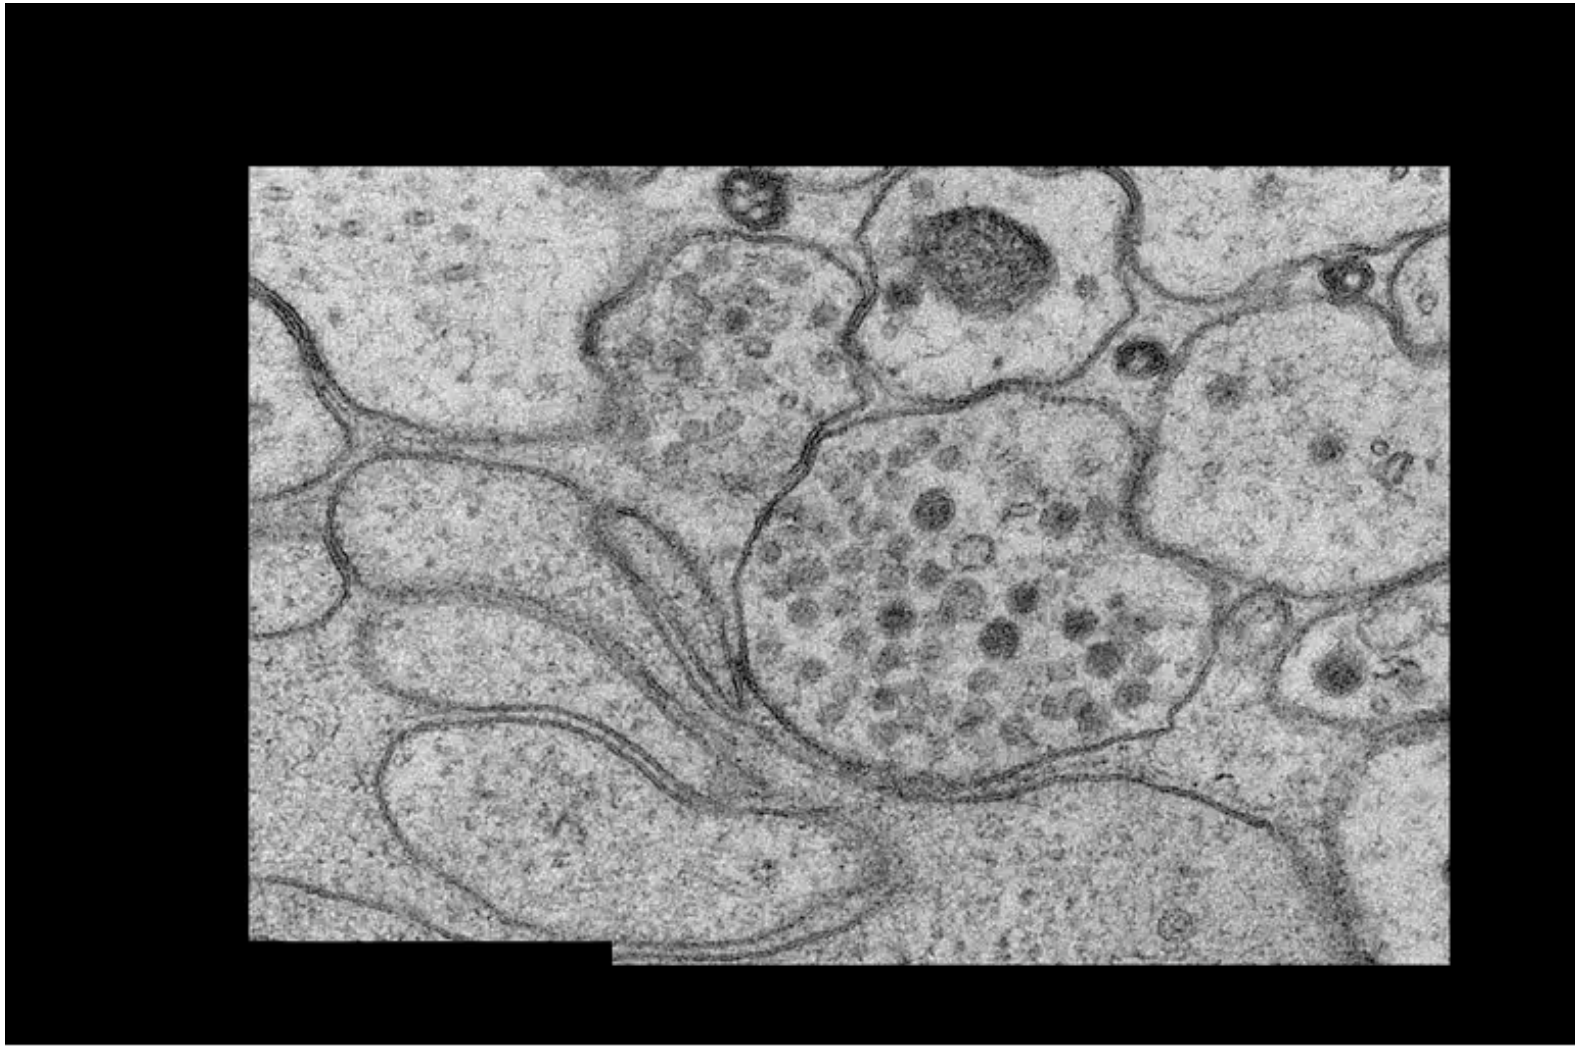

SCI INSTITUTE

C. Elegans - Jorgensen

SCI INSTITUTE

SBFSEM Images - Chien, Denk

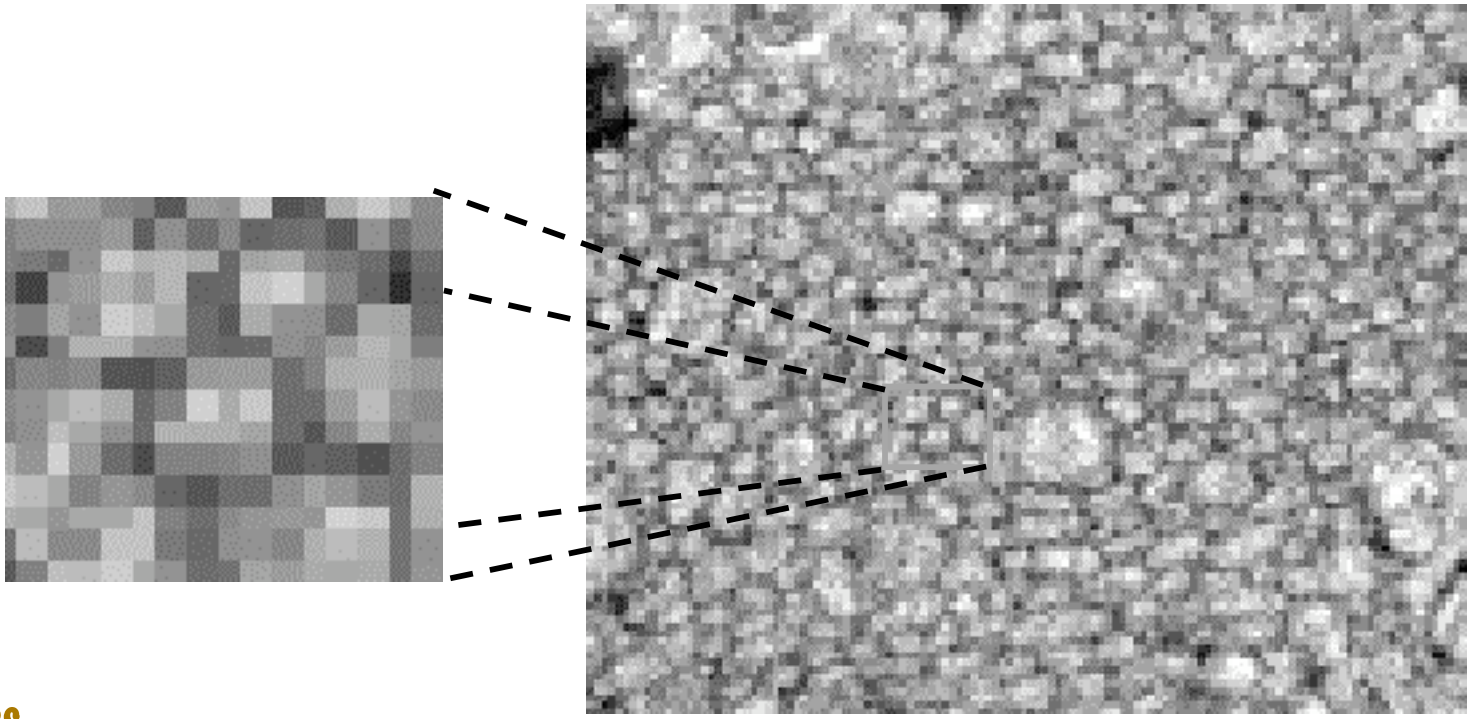

- Challenges
  - Axonal cross-sections hard to see with the eye
  - Anisotropic resolution (26x26x50nm)
  - Lower signal to noise ratio than TEM
- Prior Knowledge
  - Cutting plane nearly perpendicular to axon
  - Axons rarely branch or terminate

Tracking Overview

1. Smoothing/Noise Removal

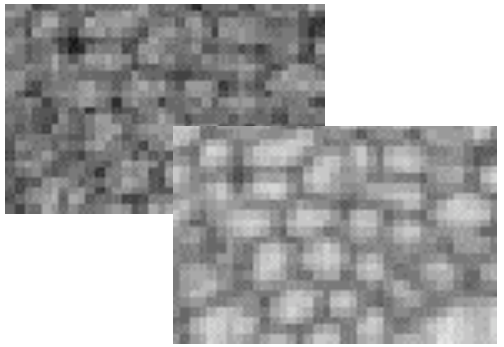

2. Axon Initialization

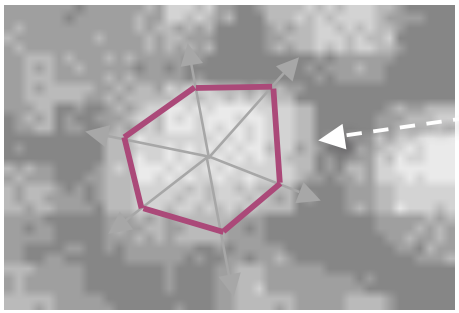

3. Axon Tracking

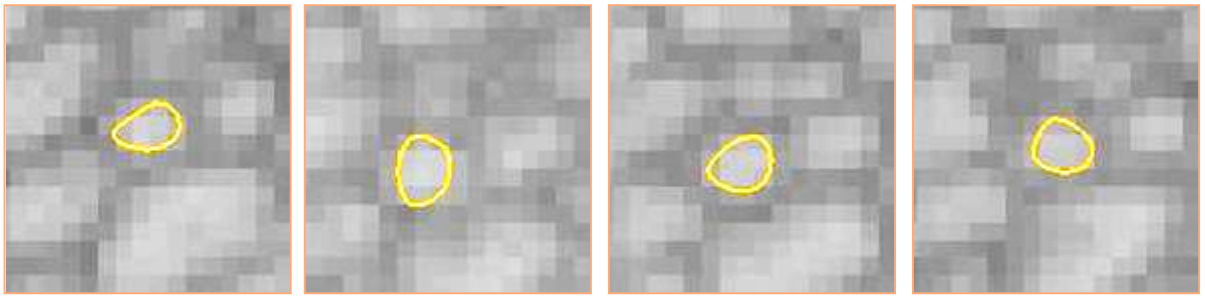

Axon Tracking

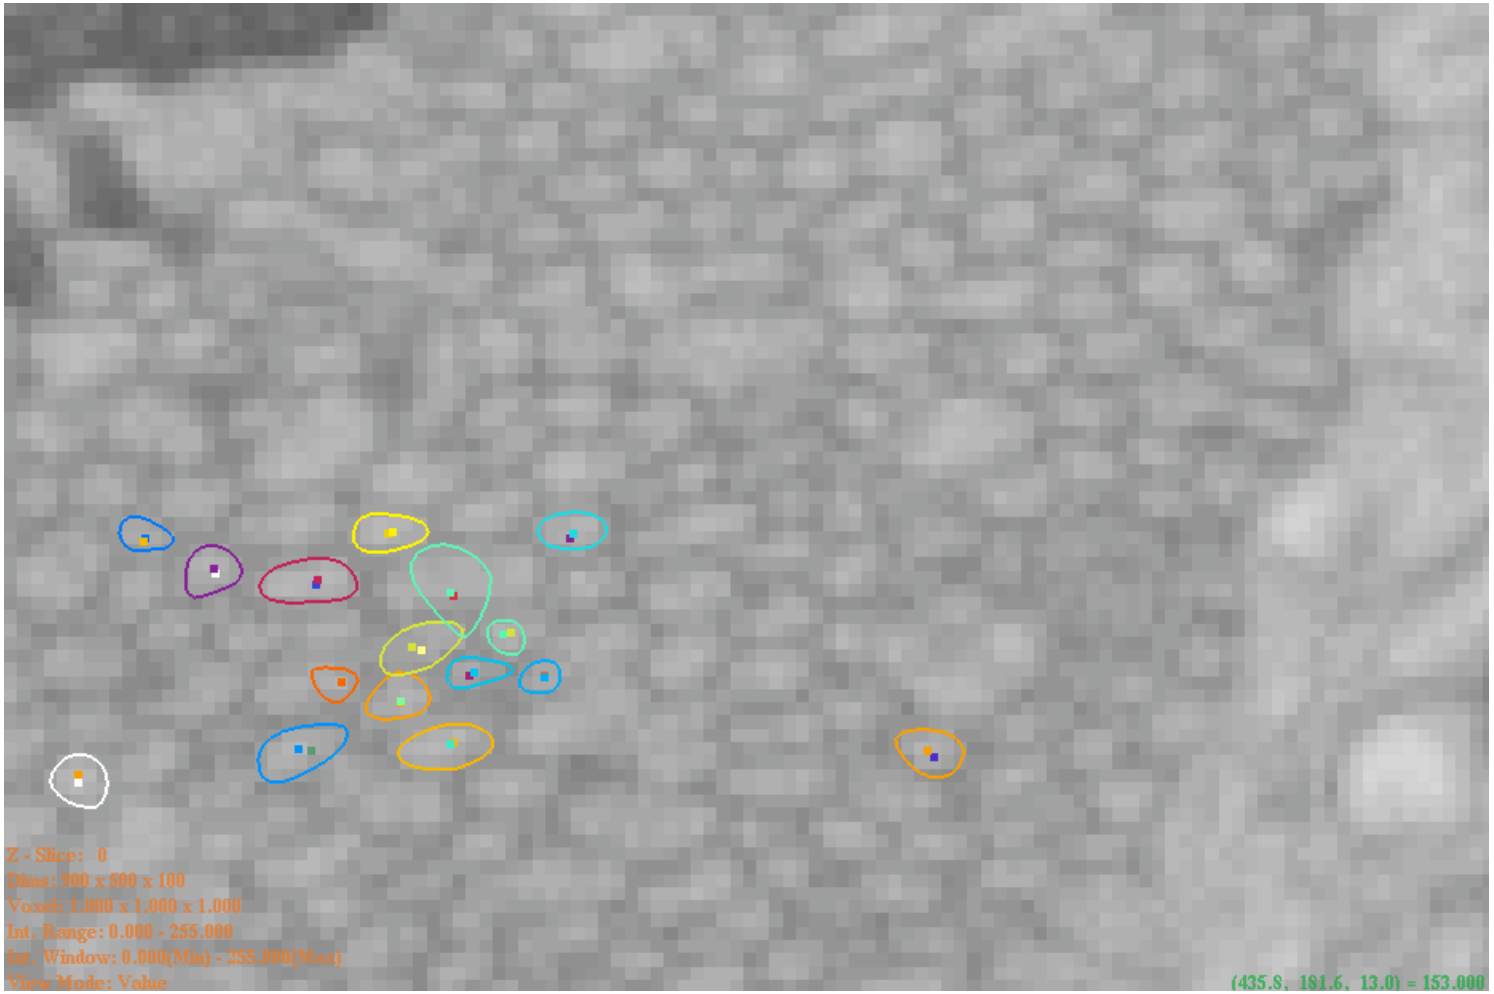

Axon Tracking

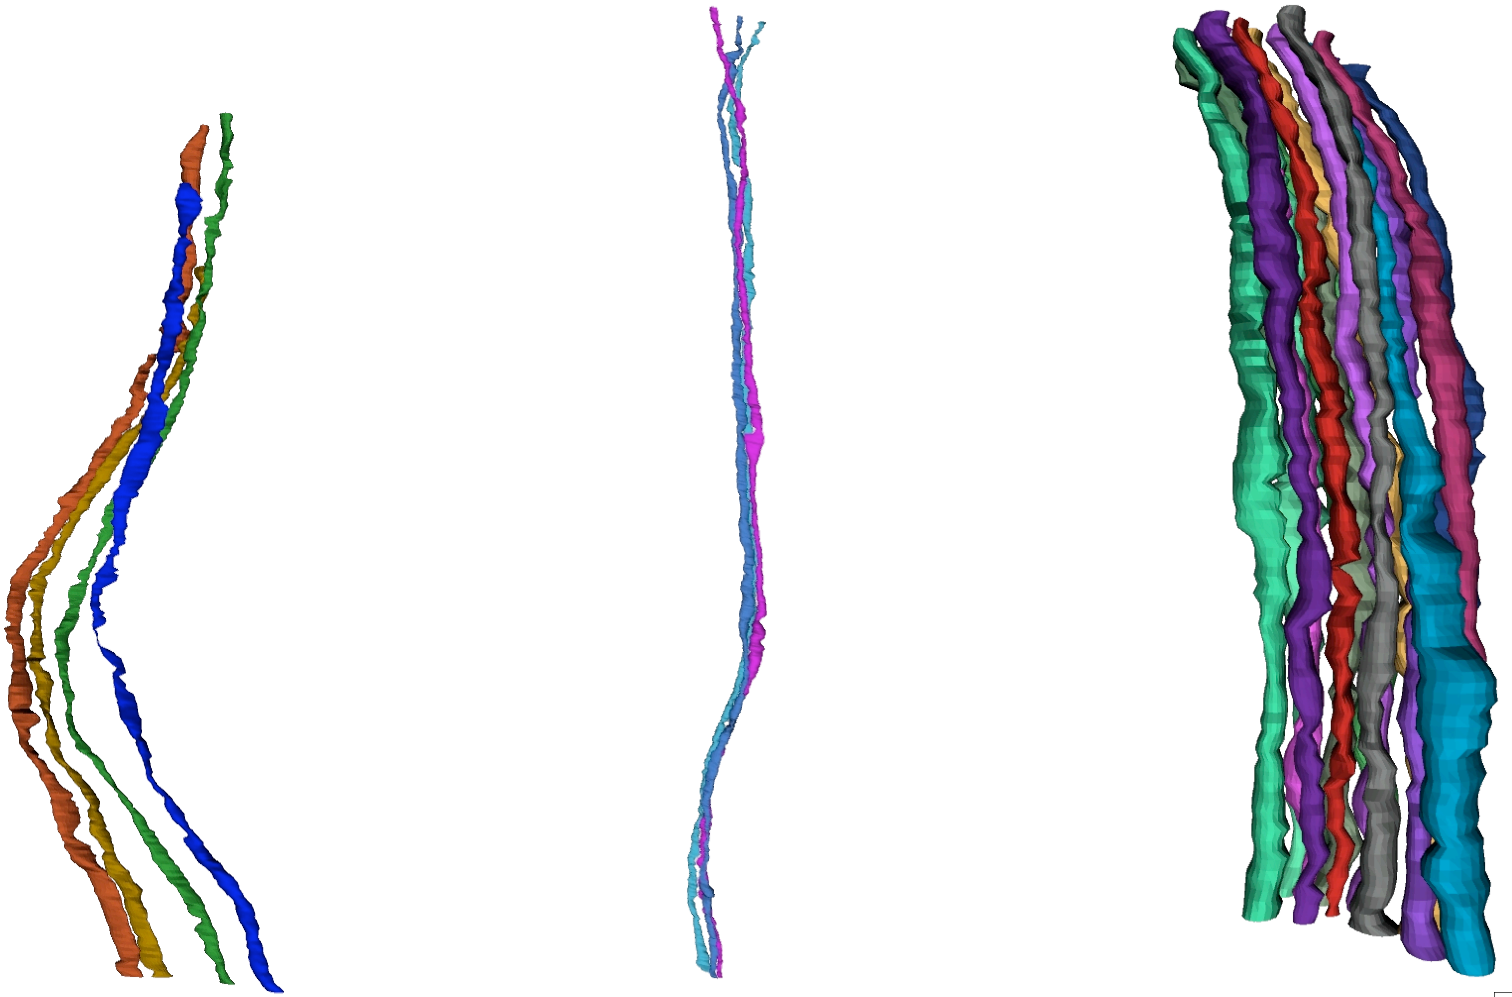

SCI INSTITUTE

Thanks

- Sponsors (NSF, NIH)
- Team: S. Awate, T. Tasdizen, N. Foster