Scalable Strategies for Image Analysis in Neuroscience

Walter J. Scheirer

Computer Vision Research Laboratory,
Department of Computer Science and Engineering

UNIVERSITY OF NOTRE DAME

Remarkable recent advances in visual recognition

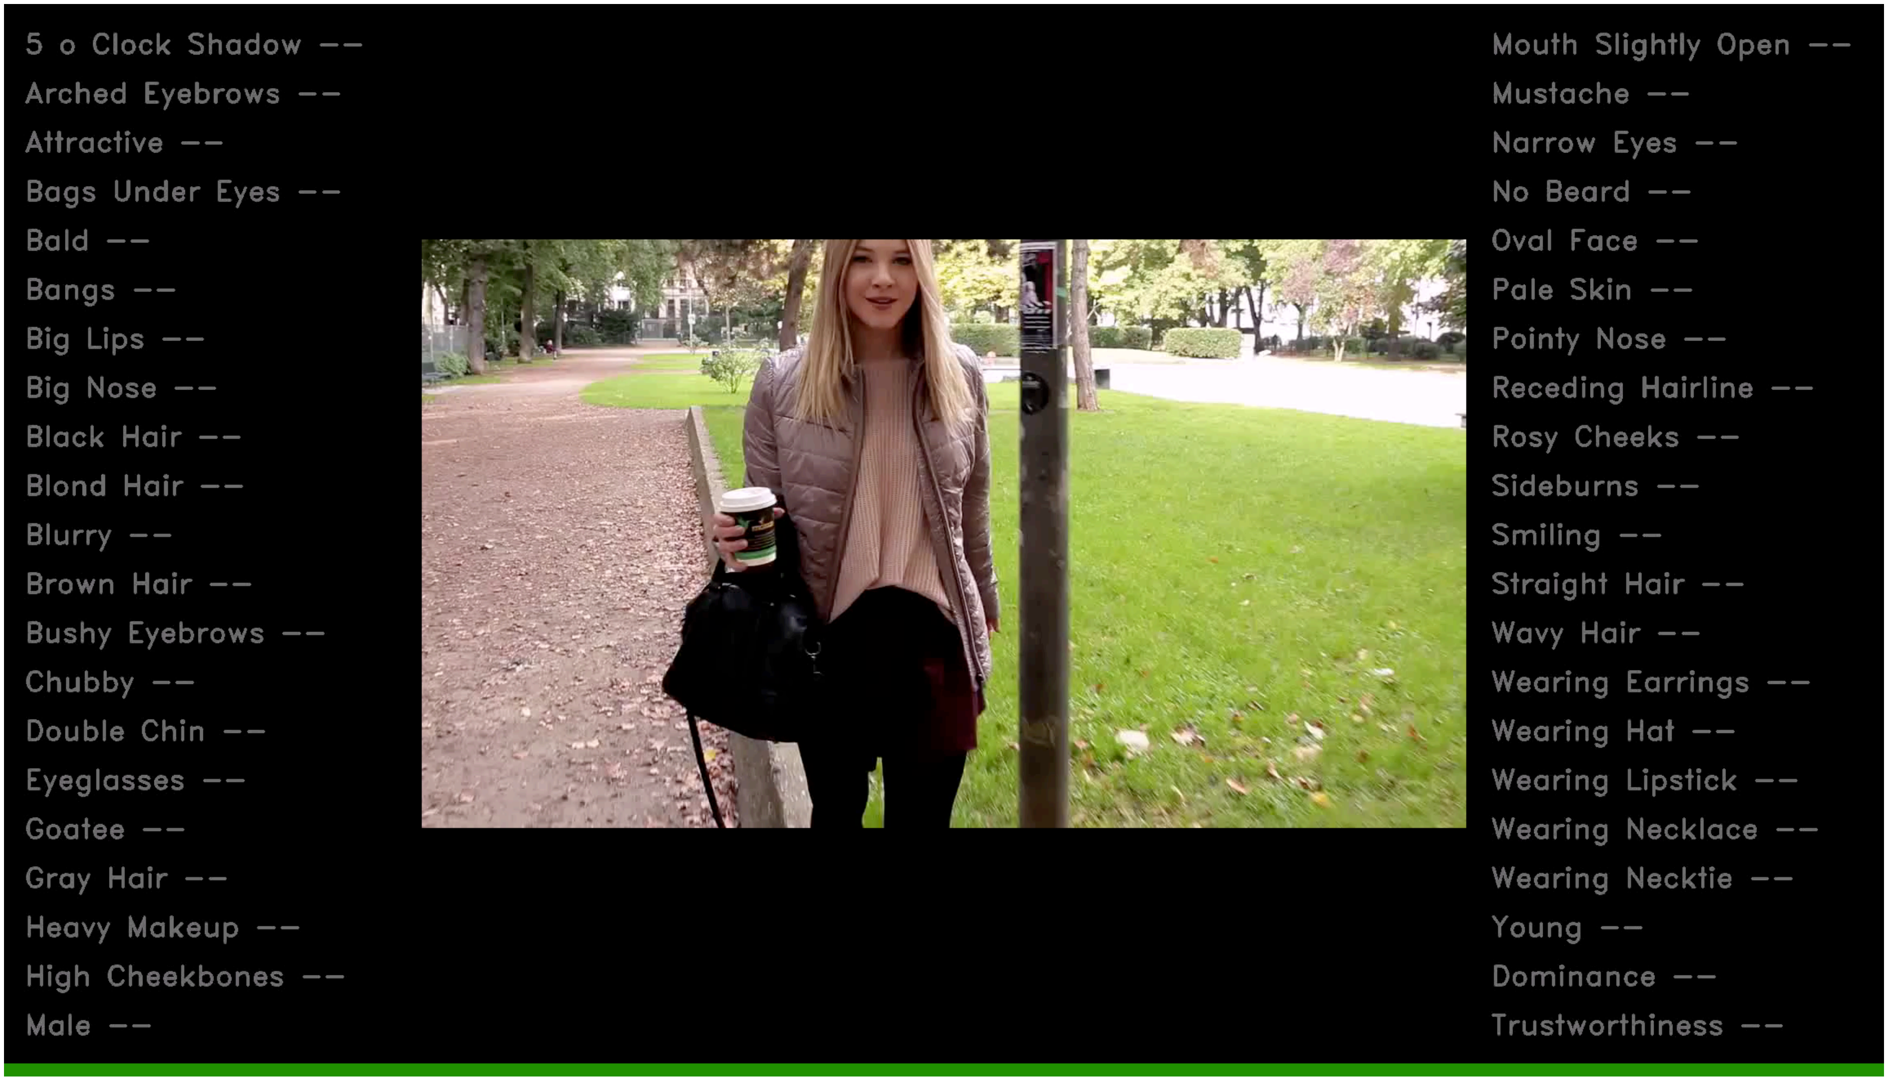

Trouble…

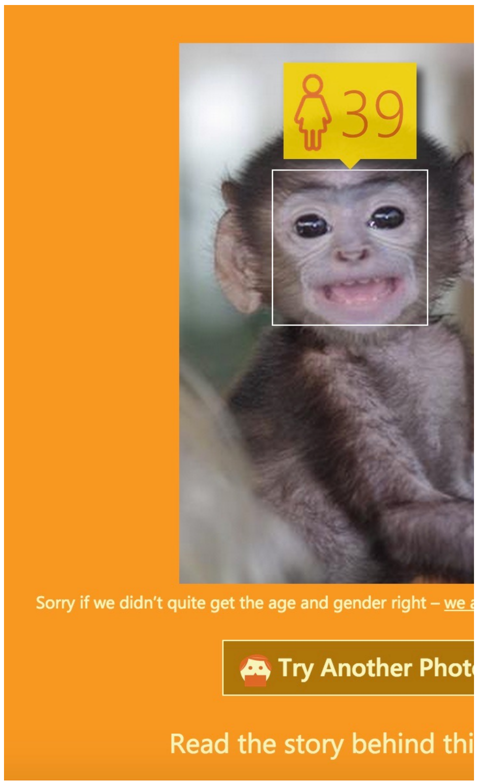

https://how-old.net/

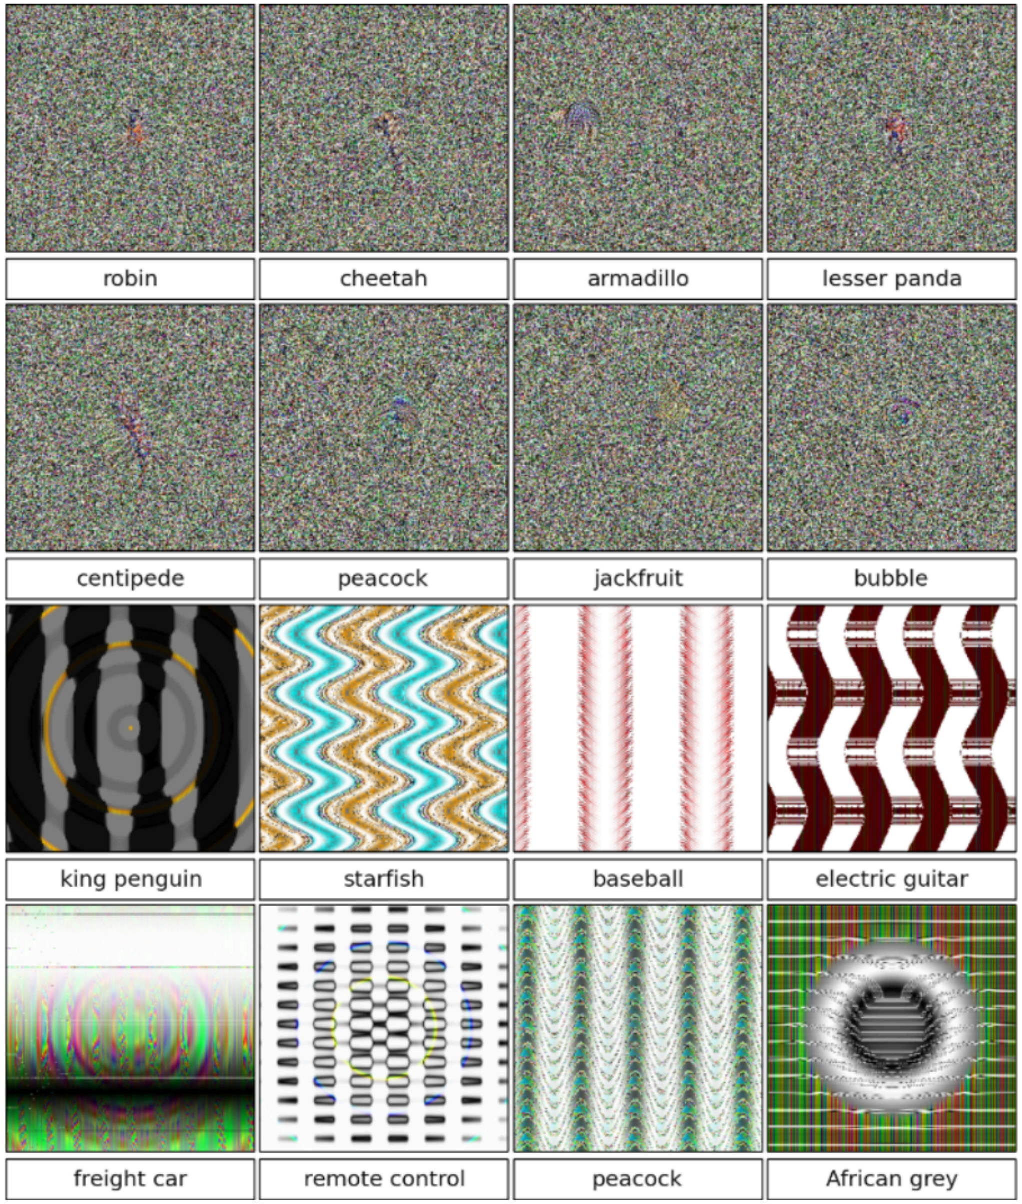

Nguyen et al. CVPR 2015

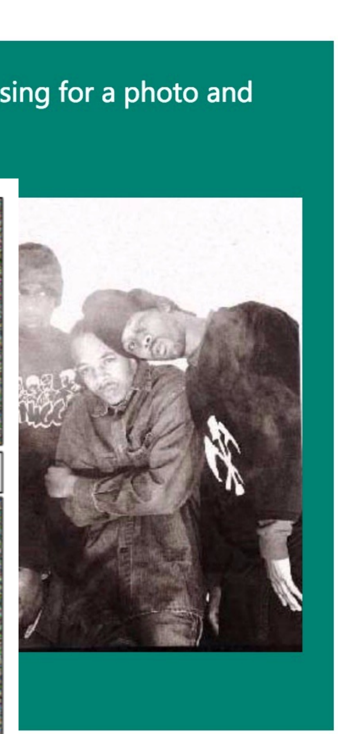

https://www.captionbot.ai/

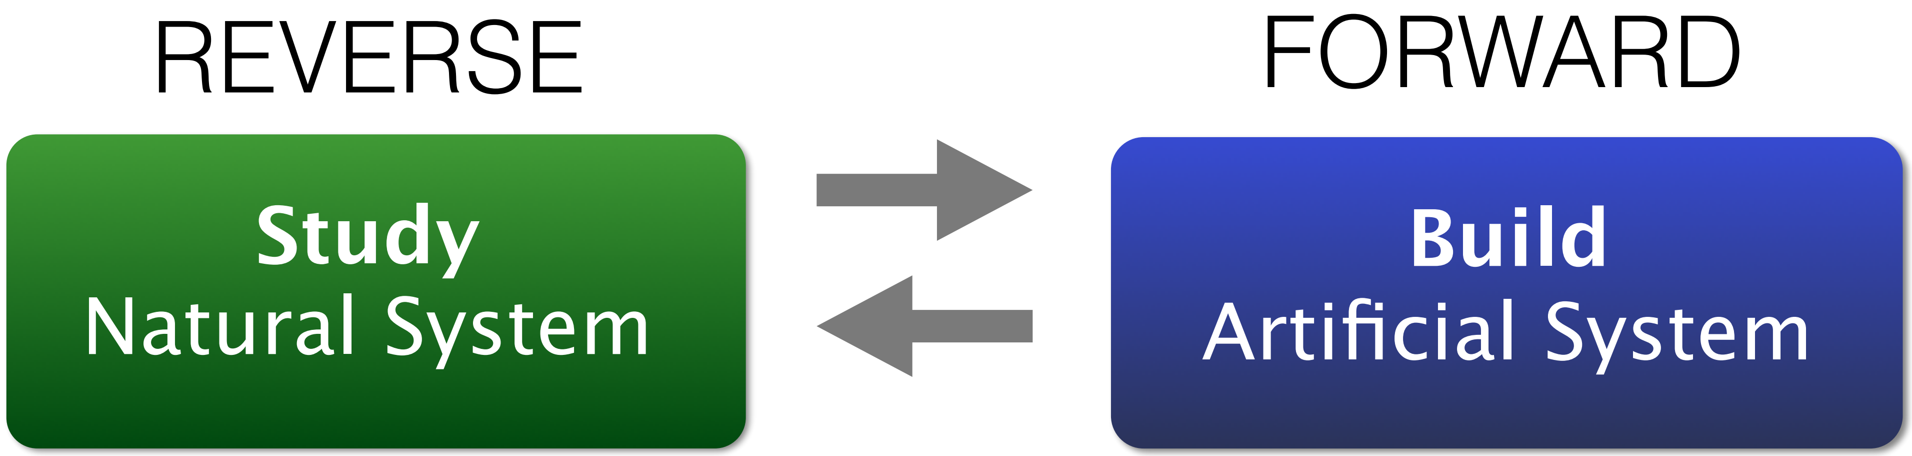

Behavior, fMRI,
2-photon Imaging,
Confocal Microscopy,
Electron Microscopy,
Electrode Recordings

Computer Vision
Machine Learning

Vision Systems that Work

Rodents

THIS

(brain)

IS FOUND IN HERE

(rattus norvegicus)

Vision Systems that Work: Models

Neuroscience and Machine Learning Have a Common Problem:

Just because a system works doesn't mean we understand it.

Even when we have considerable access to the system, it's not always easy to answer "why" questions

Maybe we can build tools that help address both issues and learn something in the process

New Experiments

Establishing causal links between cortical areas and function

Understanding decision boundaries between classes

Studying the mechanisms of sensory integration

Searching for mesoscale cortical computing circuits

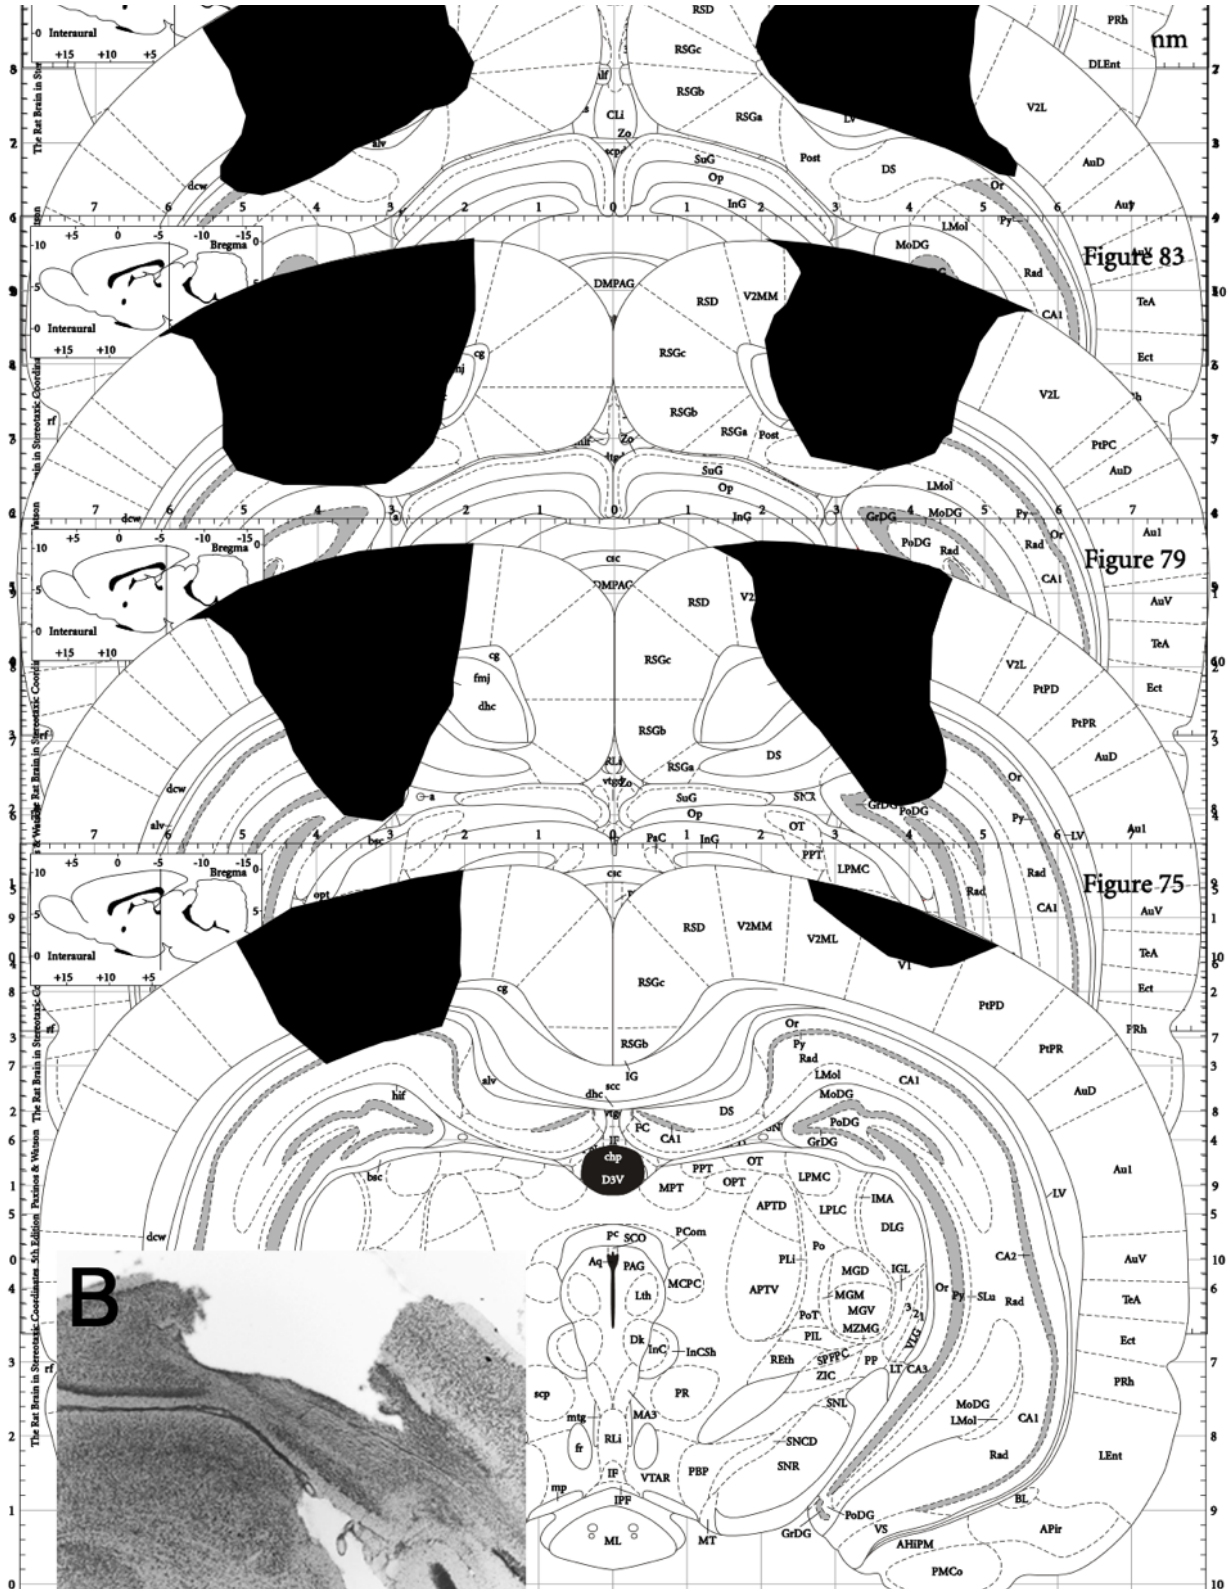

Brain Science at the Notre Dame Computer Vision Research Lab

IARPA MICrONS

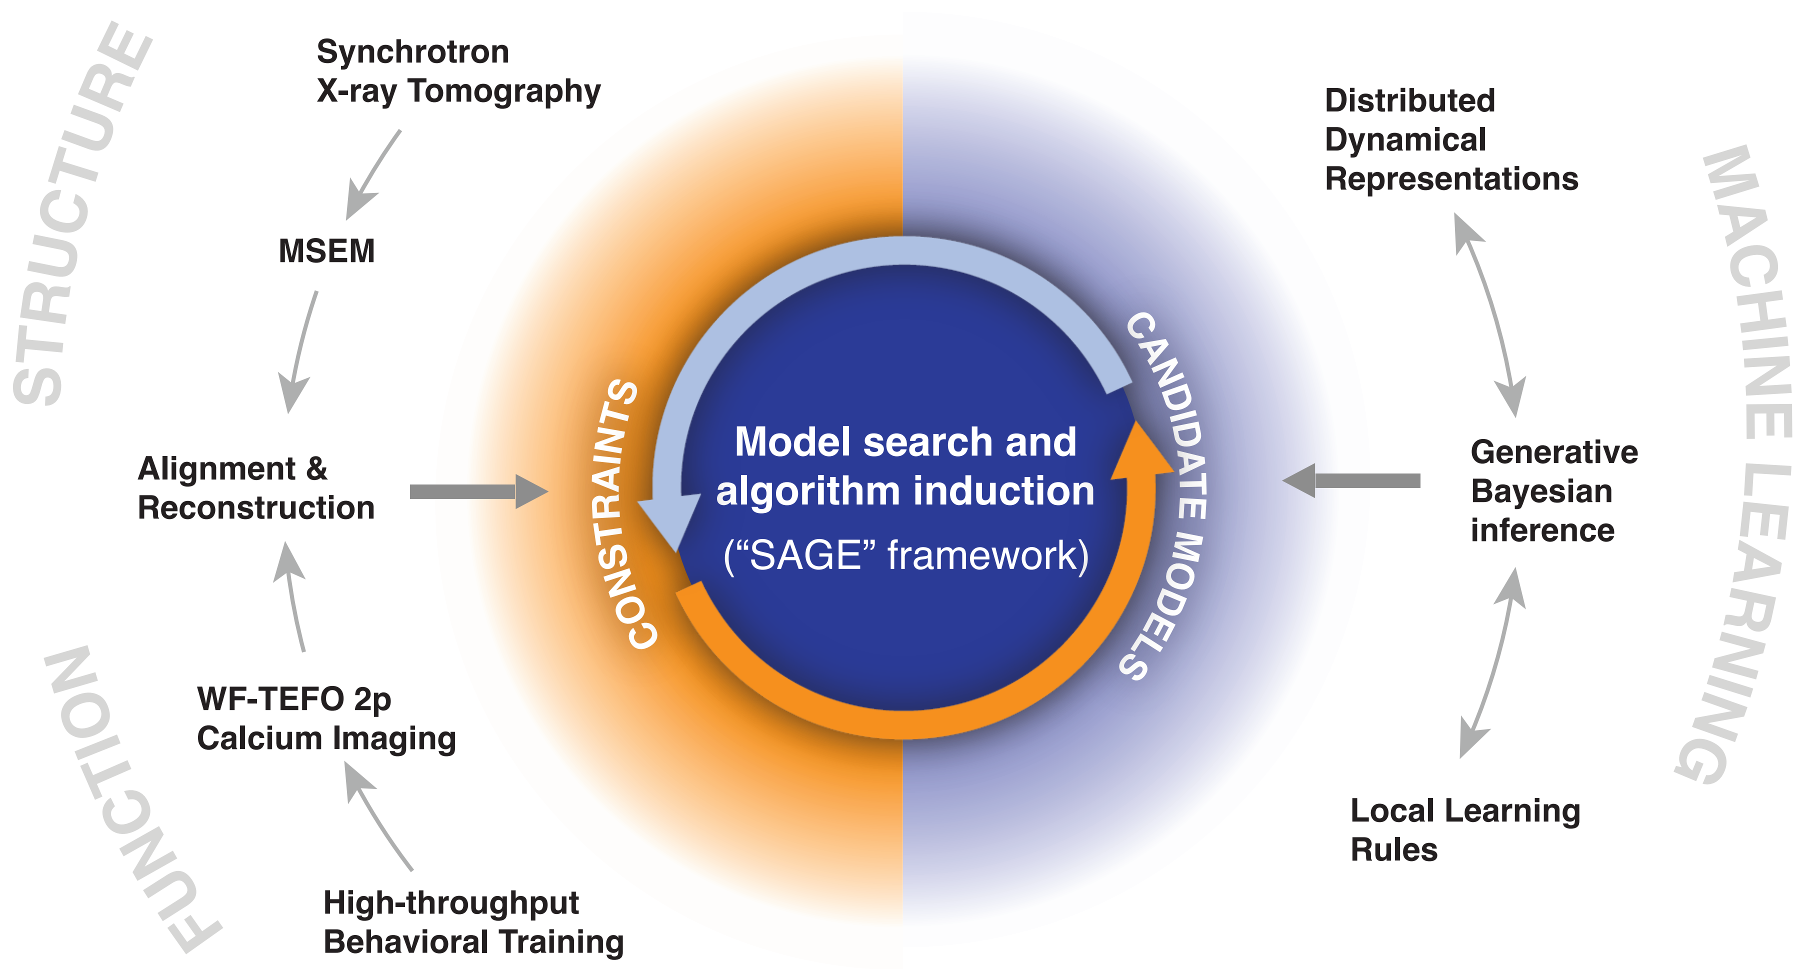

Machine Learning Algorithms from Wet Lab Experimentation

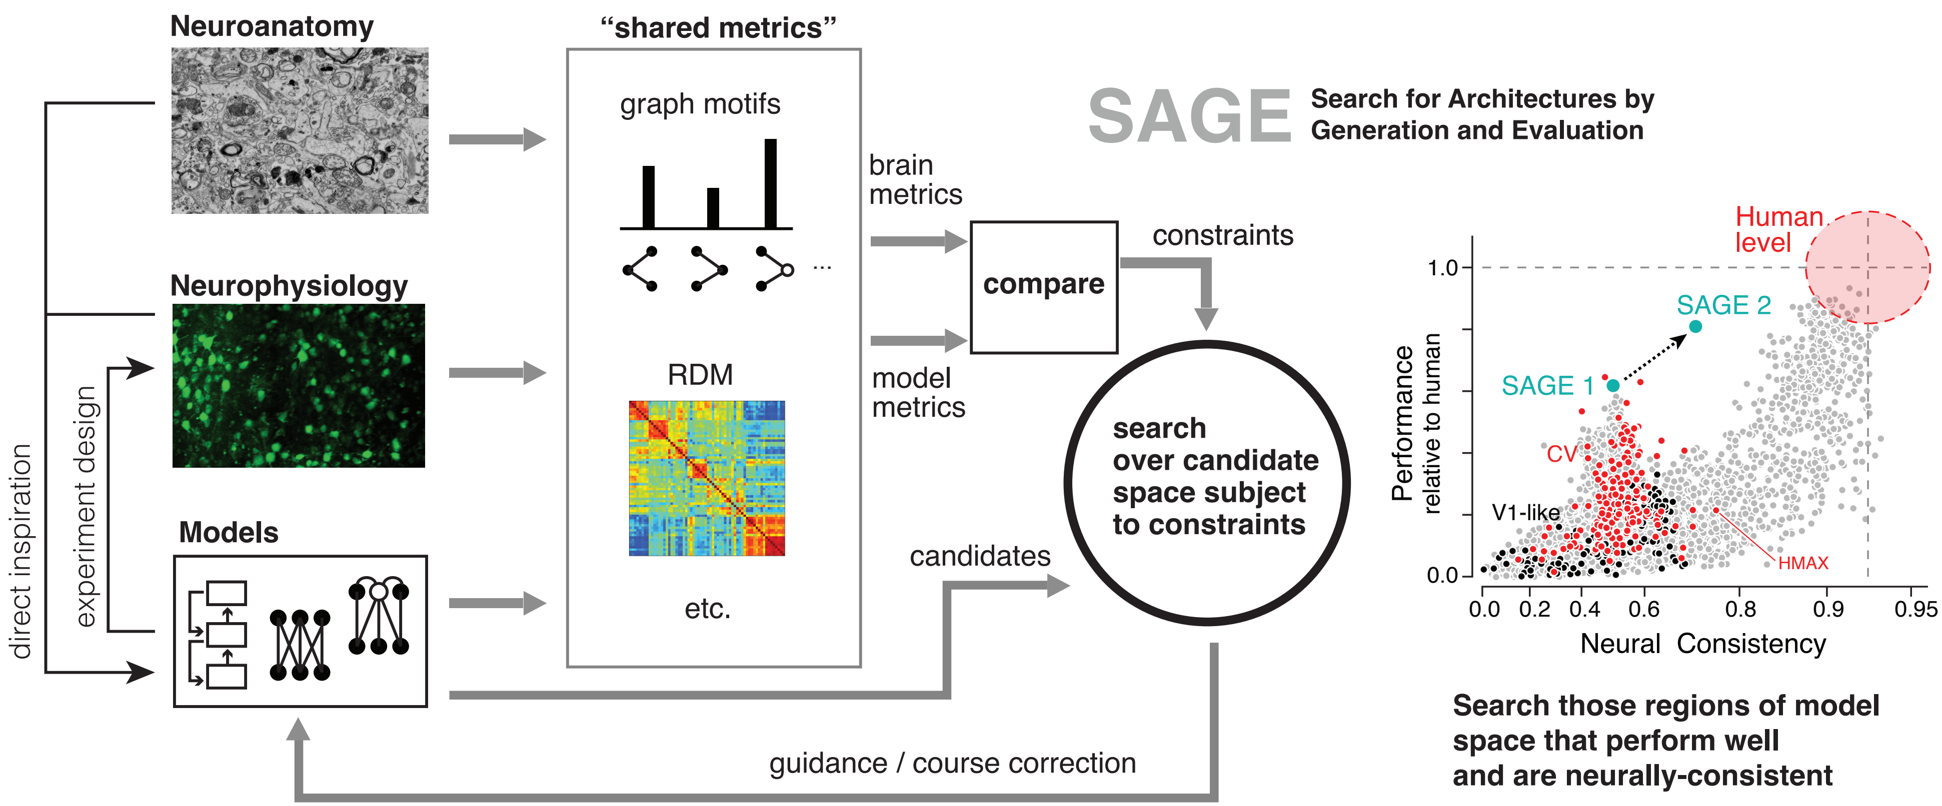

Search those regions of model space that perform well and are neurally-consistent

Experimental workflow

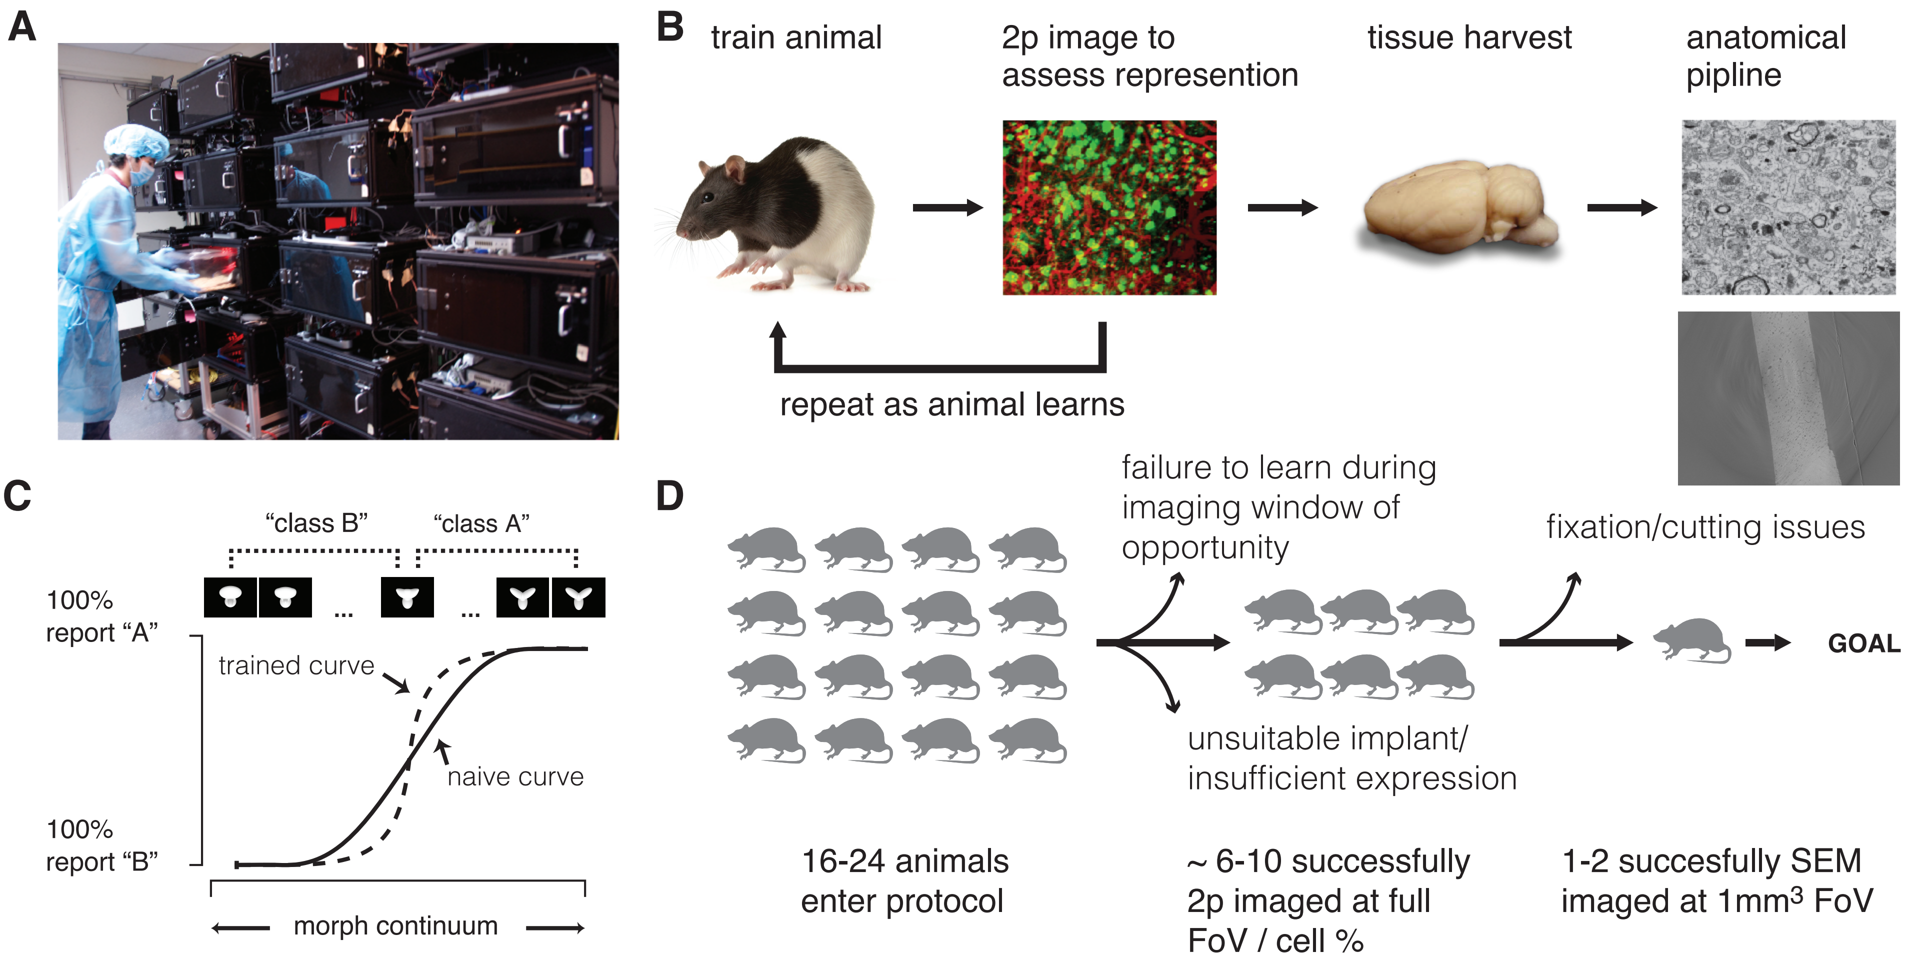

2D Segmentation and 3D Reconstruction for EM

Elia Shabazi

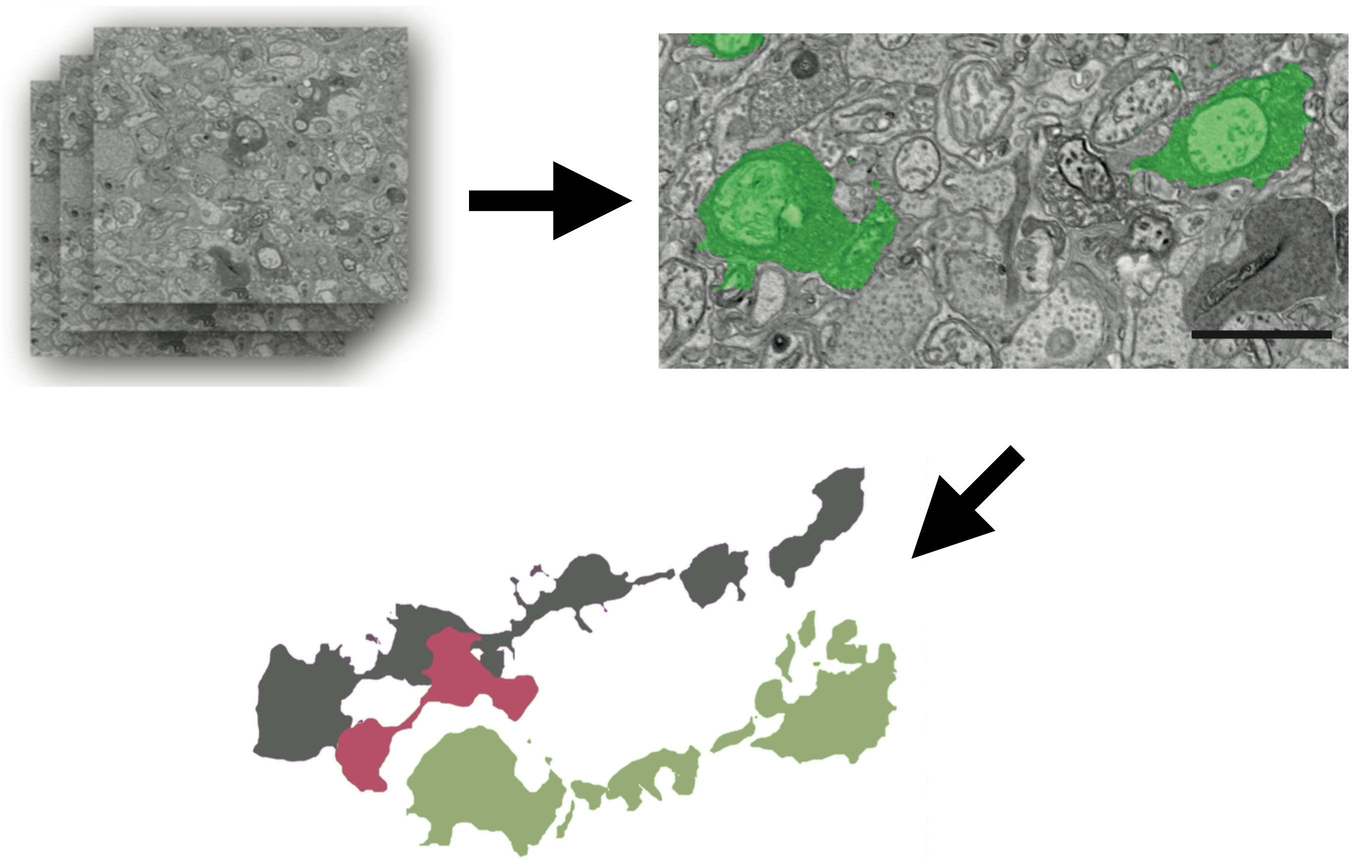

EM: M. Joesch, Cox Lab @ Harvard

2D Segmentation and 3D Reconstruction for X-ray

Elia Shabazi

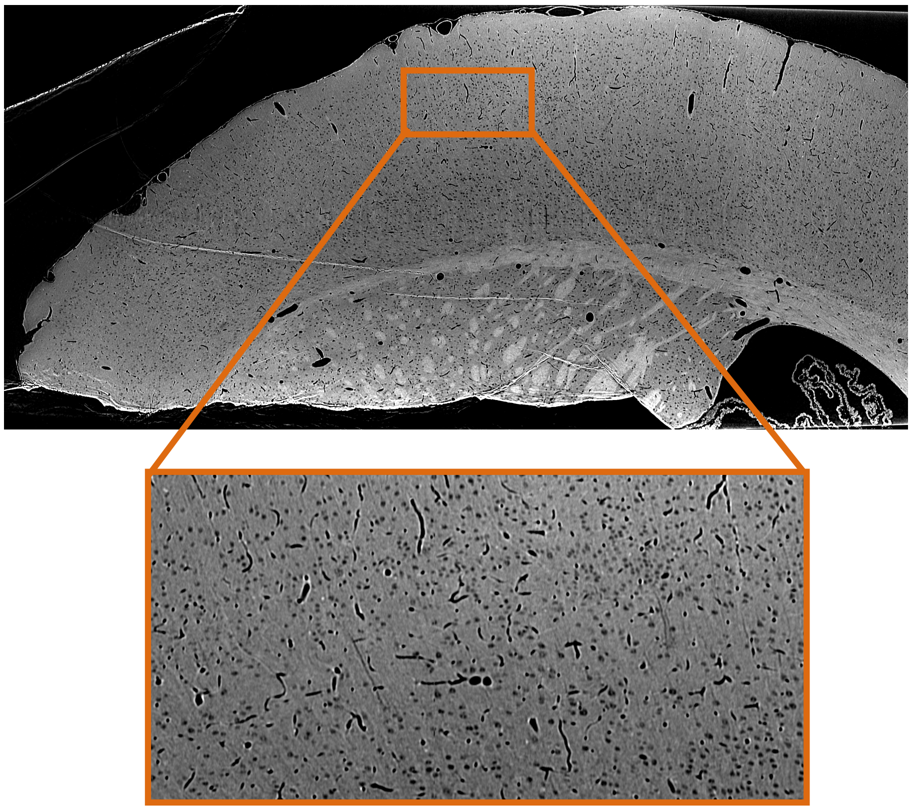

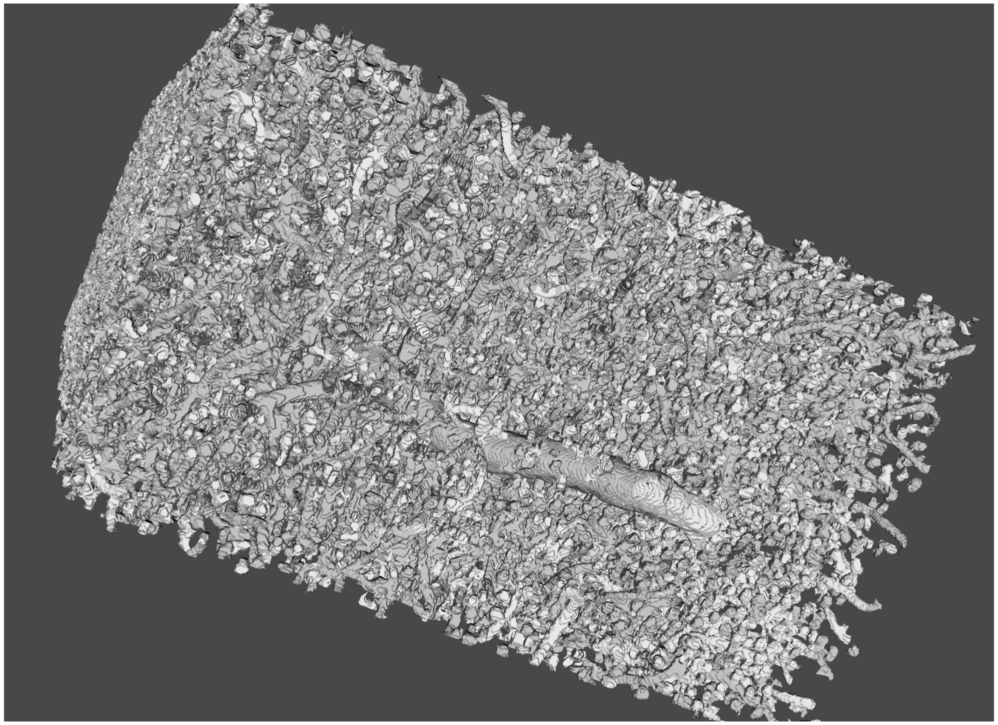

X-Ray: N. Kasthuri, Argonne National Laboratory / U. Chicago

Psychophysics on the Model

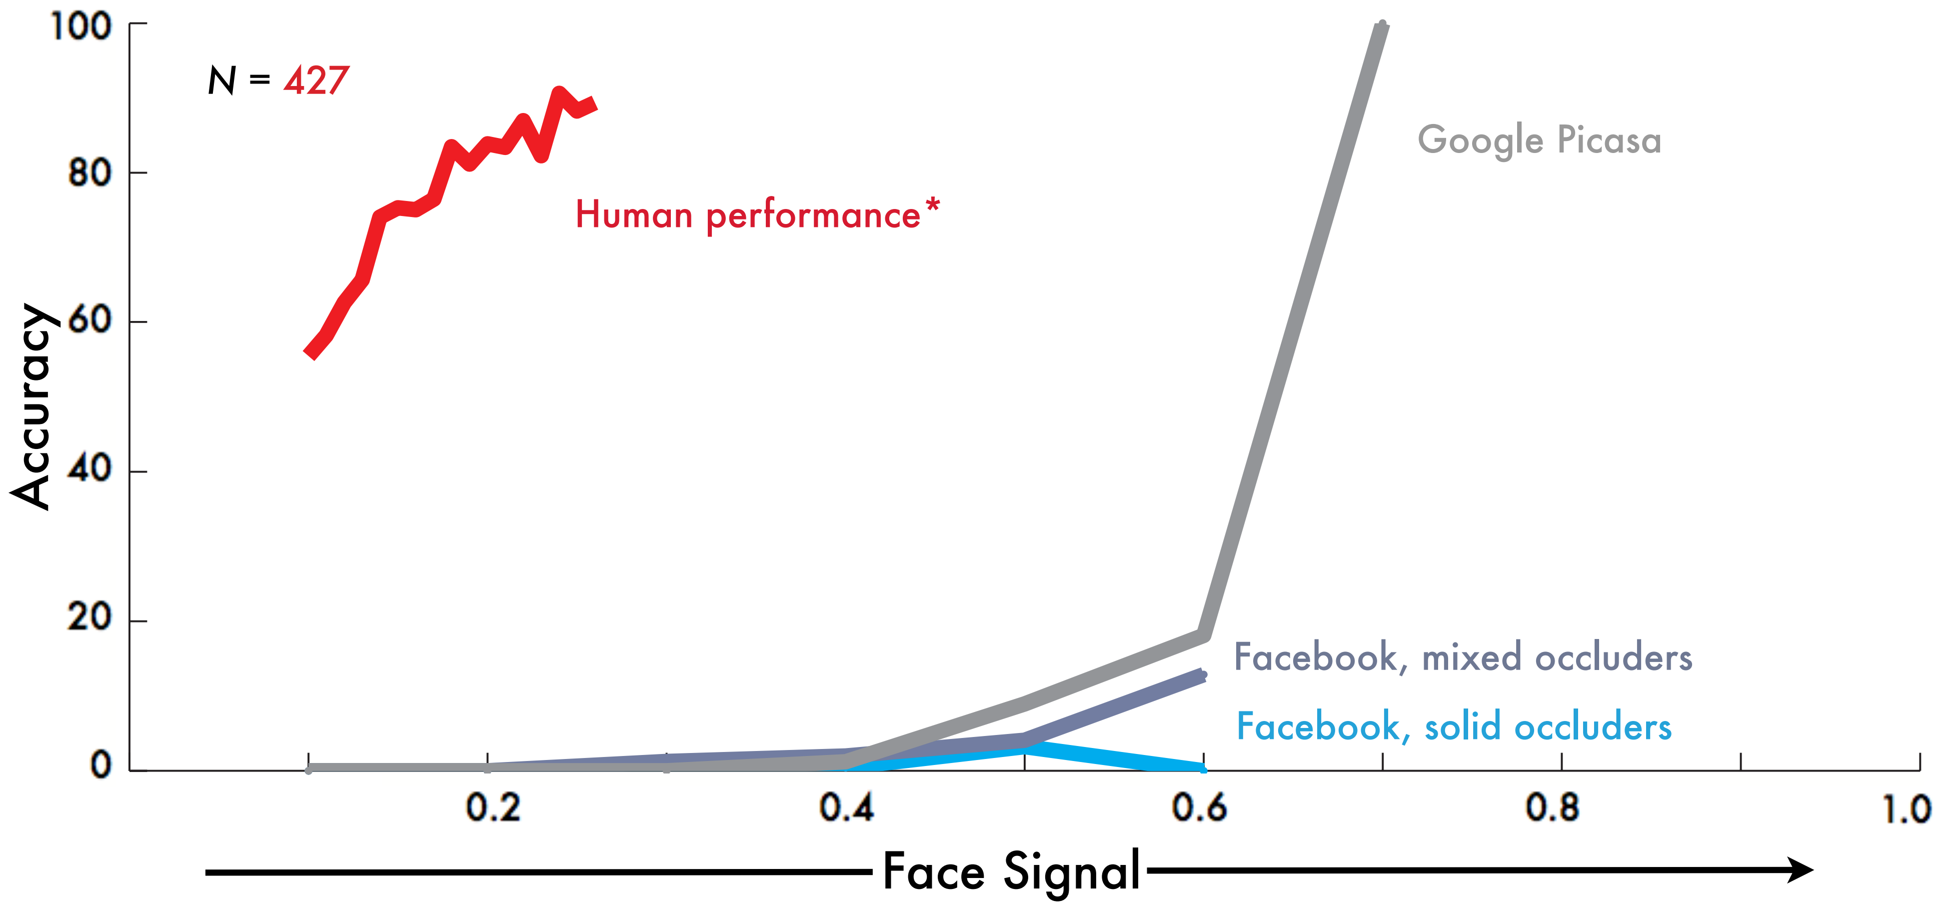

Brandon RichardWebster

Strategies for Image Analysis

Dense Segmentation and Reconstruction

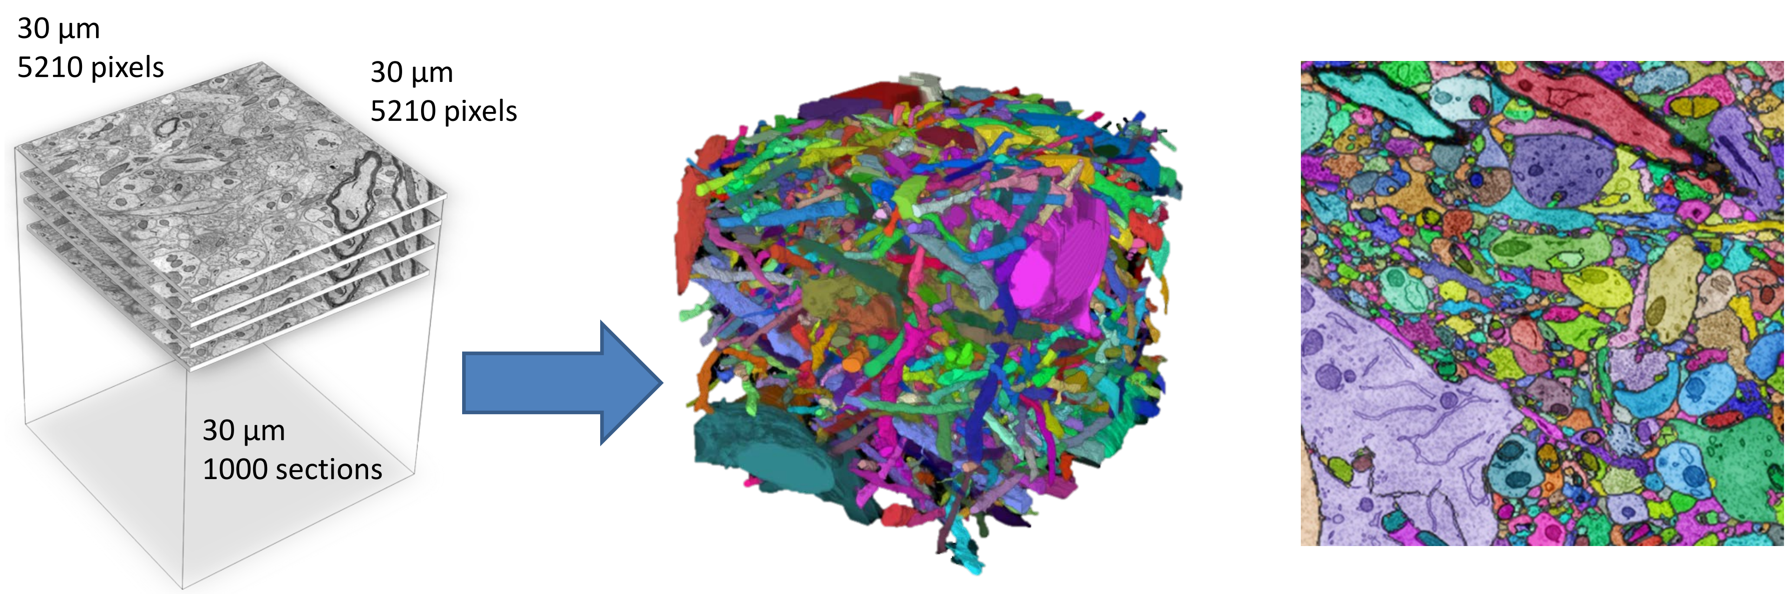

Visual Computing Group @ Harvard

Software tools for dense reconstruction

Package	**Method**	**Link**
Randomer Forests	Decision Forests	http://ttomita.github.io/RandomerForest
Gala	Active Learning	https://github.com/janelia-flyem/gala
VESICLE	Deep Learning	http://openconnecto.me/vesicle
Synapse Segmenter	Context Features + Adaboost	http://cvlab.epfl.ch/software/synapse
ATMA	3D Pixel Features + Random Forests	https://github.com/RWalecki/ATMA
ZNN	CNN	https://github.com/seung-lab/znn-release
PRIM	CRF	http://github.com/funkey/prim
ilastik	Random Forests	http://ilastik.org/
Rhoana	CNN	https://github.com/Rhoana

Preeminent 2D segmentation method: CNN

Feature learning for strongly invariant membrane representations

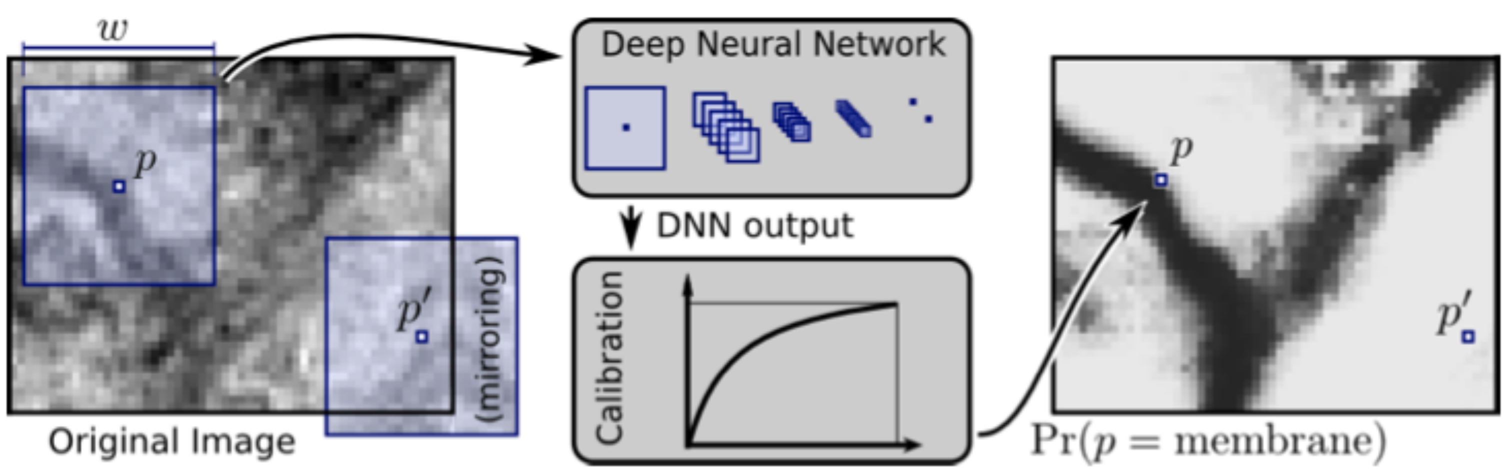

Cirşan et al. NIPS 2012

Interlude: is deep learning as good as we think it is?

An alternative to dataset testing

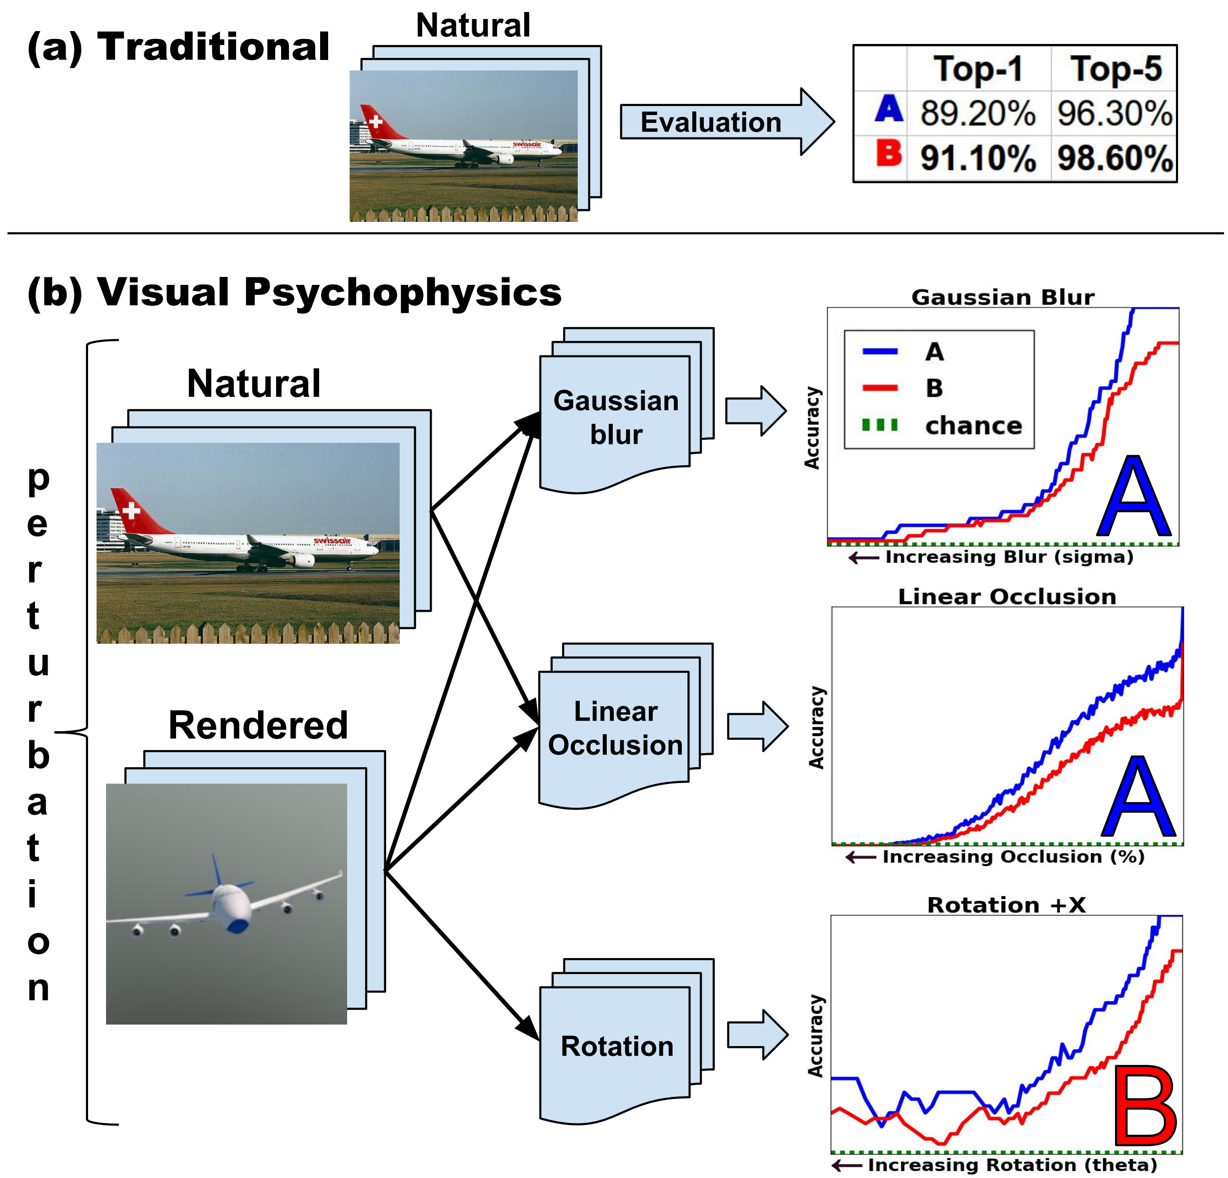

B. RichardWebster, S. Anthony, and W. J. Scheirer, "PsyPhy: A Psychophysics Driven Evaluation Framework for Visual Recognition," arXiv:1611.06448, 2016

The framework

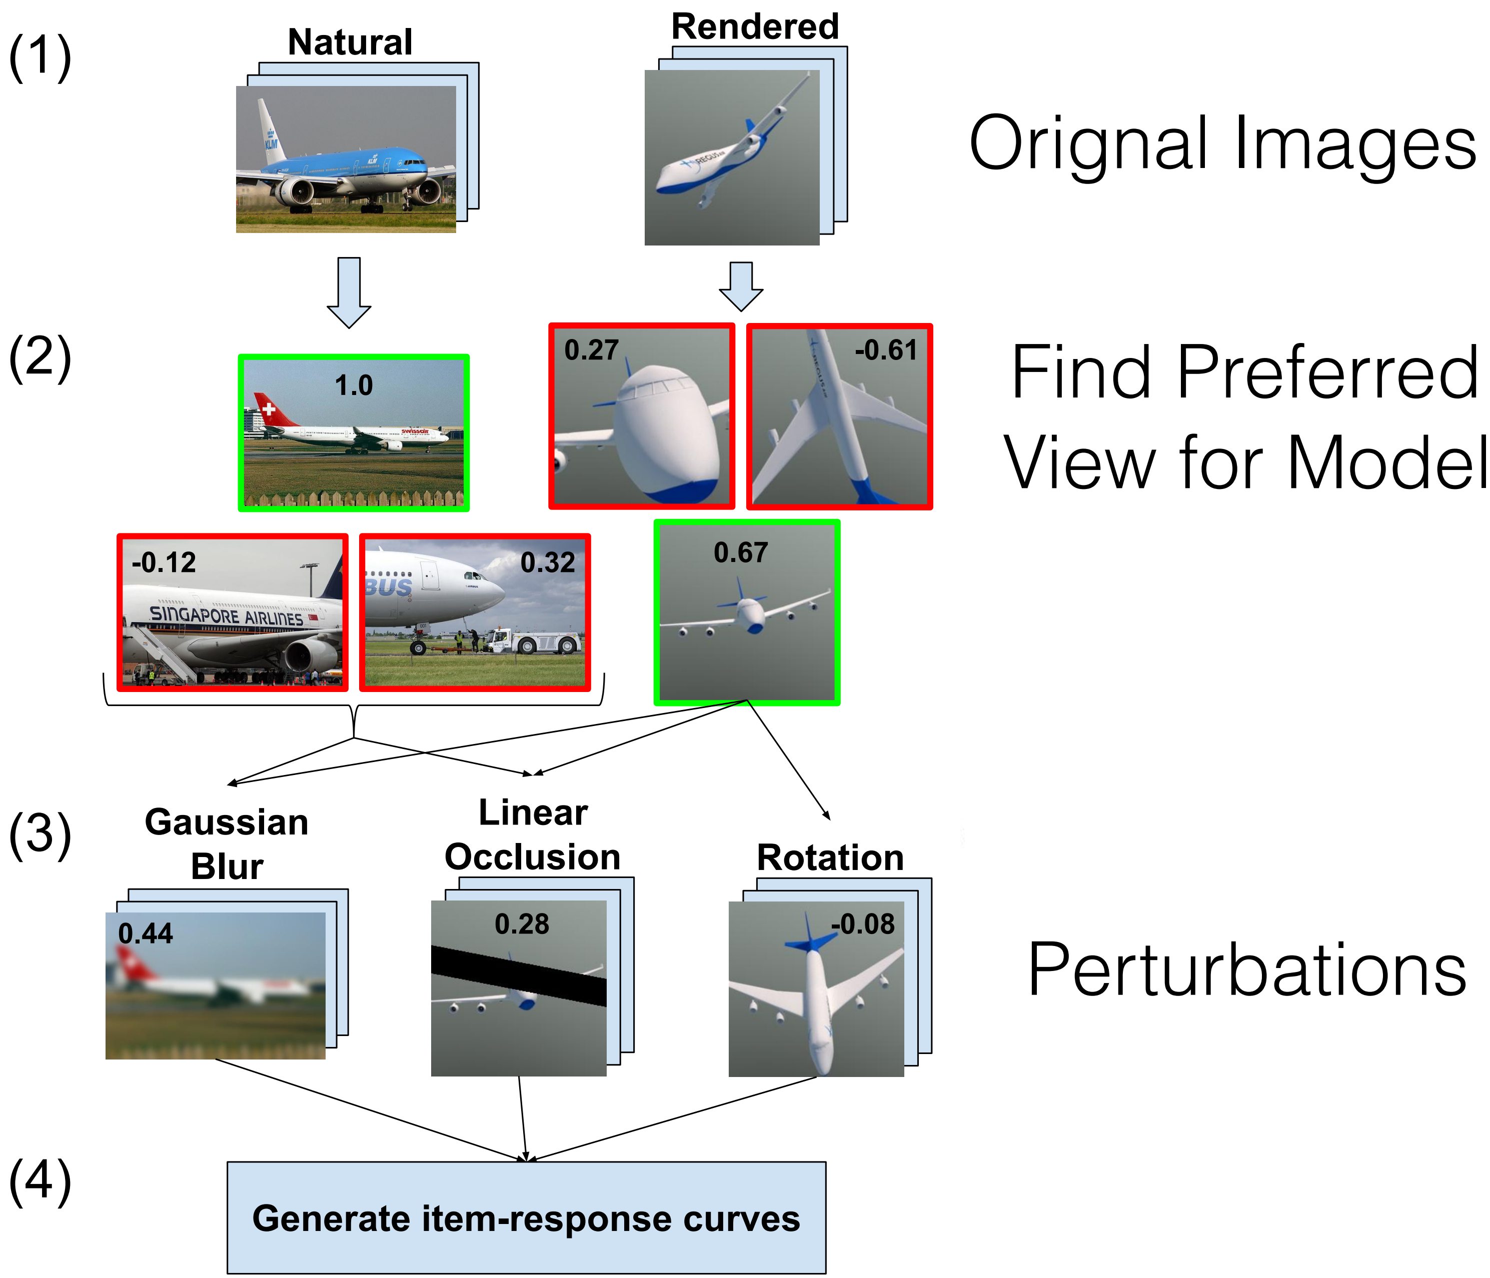

You don’t have to use tricky manipulations

GoogleNet Output

Label: Hammerhead Shark

Label: Blow Dryer

Label: Mosque

Label: Syringe

Label: Trimaran

Label: Missile

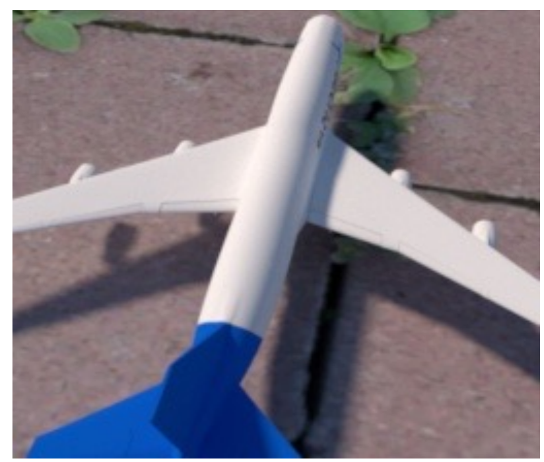

Same transformation across 2D and 3D objects

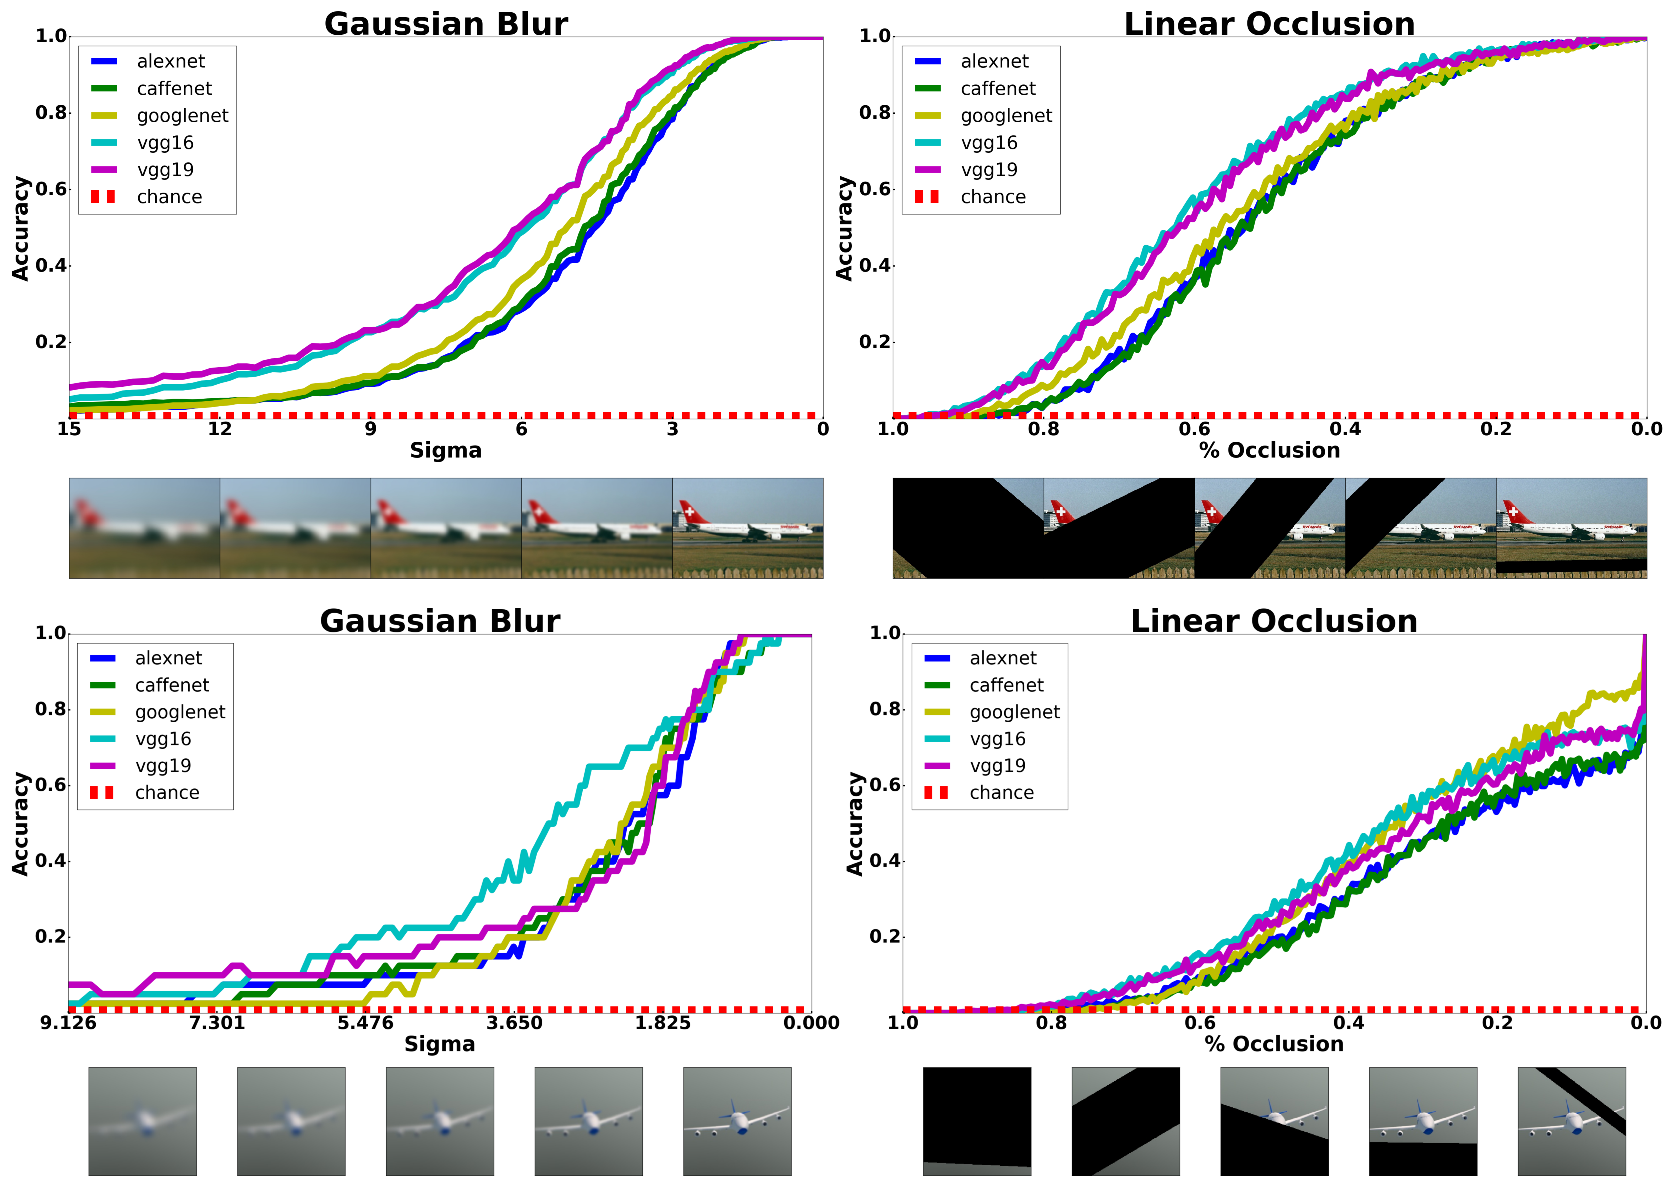

Various transformations across 3D objects

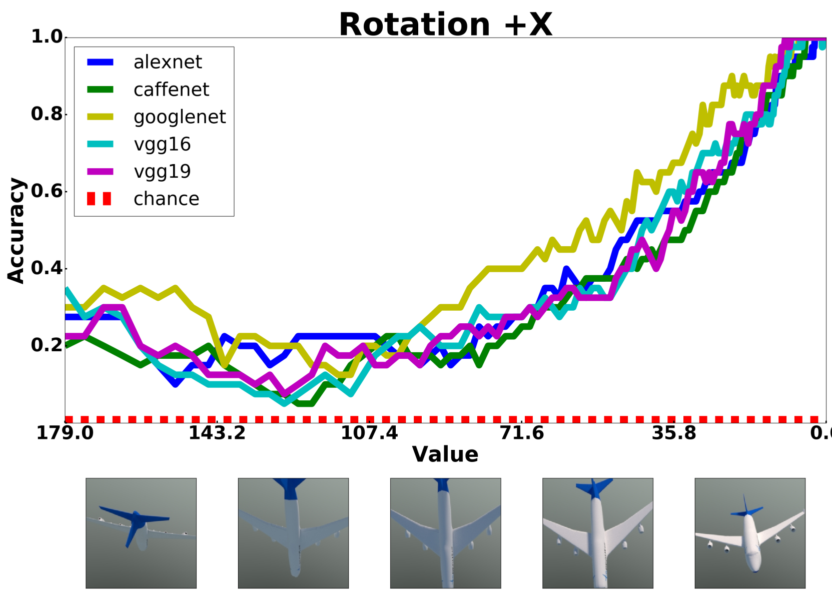

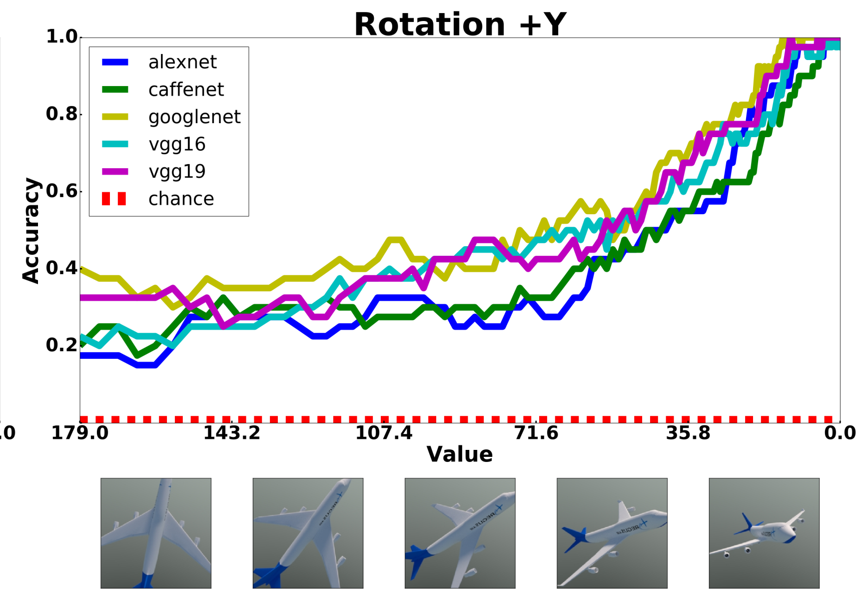

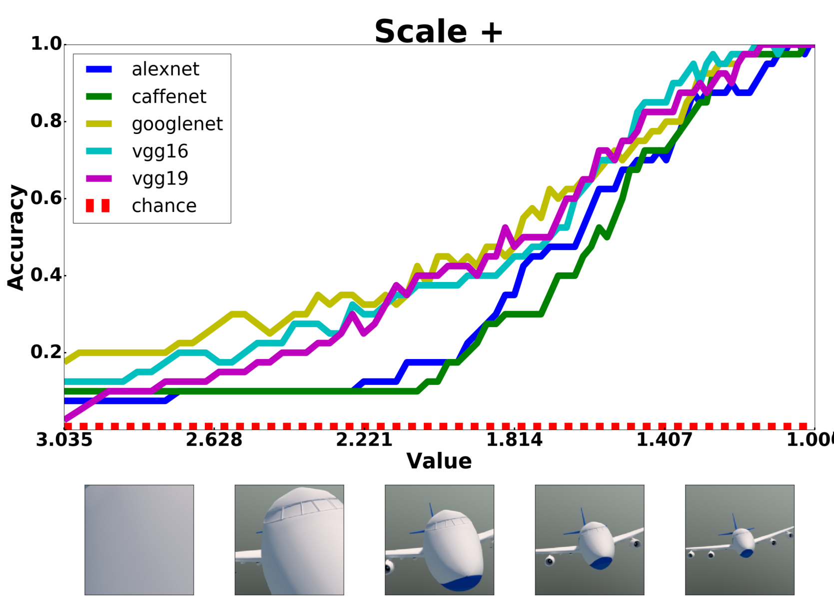

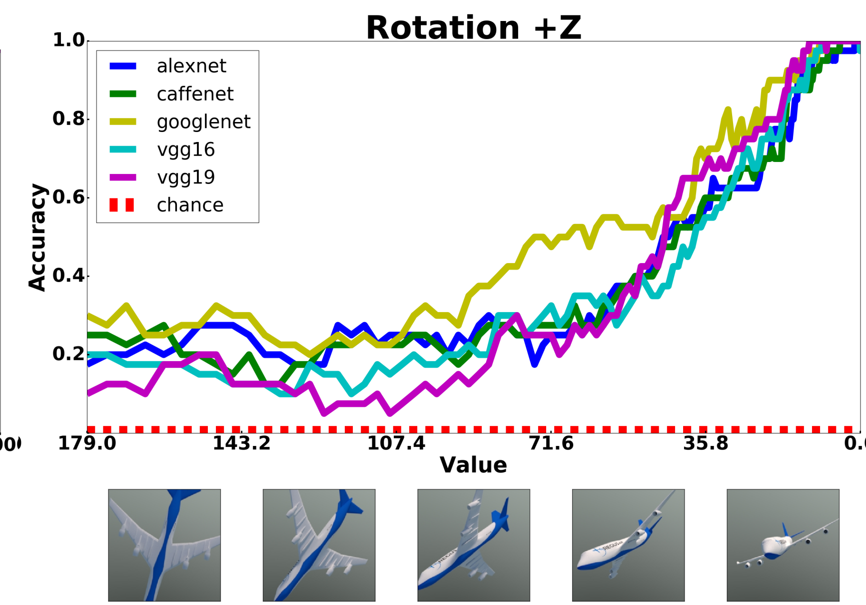

Back to image analysis…

Error rates (MICrONS targets)

100 x 100 x 100 μm^3	**1mm x 1mm x 0.1mm**	**1mm x 1mm x 1mm**
Precision: ≥ 70%	Precision: ≥ 85%	Precision: ≥ 97.5%
Recall: ≥ 70%	Recall: ≥ 85%	Recall: ≥ 97.5%
NID: ≤ 0.95 @ 50th pctl.	NID: ≤ 0.95 @ 15th pctl.	NID: ≤ 0.80 @ 10th pctl.
≤ 0.50 @ 75th pctl.	≤ 0.35 @ 50th pctl.	≤ 0.15 @ 50th pctl.
≤ 0.20 @ 95th pctl.	≤ 0.15 @ 85th pctl.	≤ 0.05 @ 75th pctl.
VI ≤ 1.75 nats	VI ≤ 1.0 nats	VI ≤ 0.25 nats

Think about error propagation for even the best of these numbers…

Training Time

AlexNet:

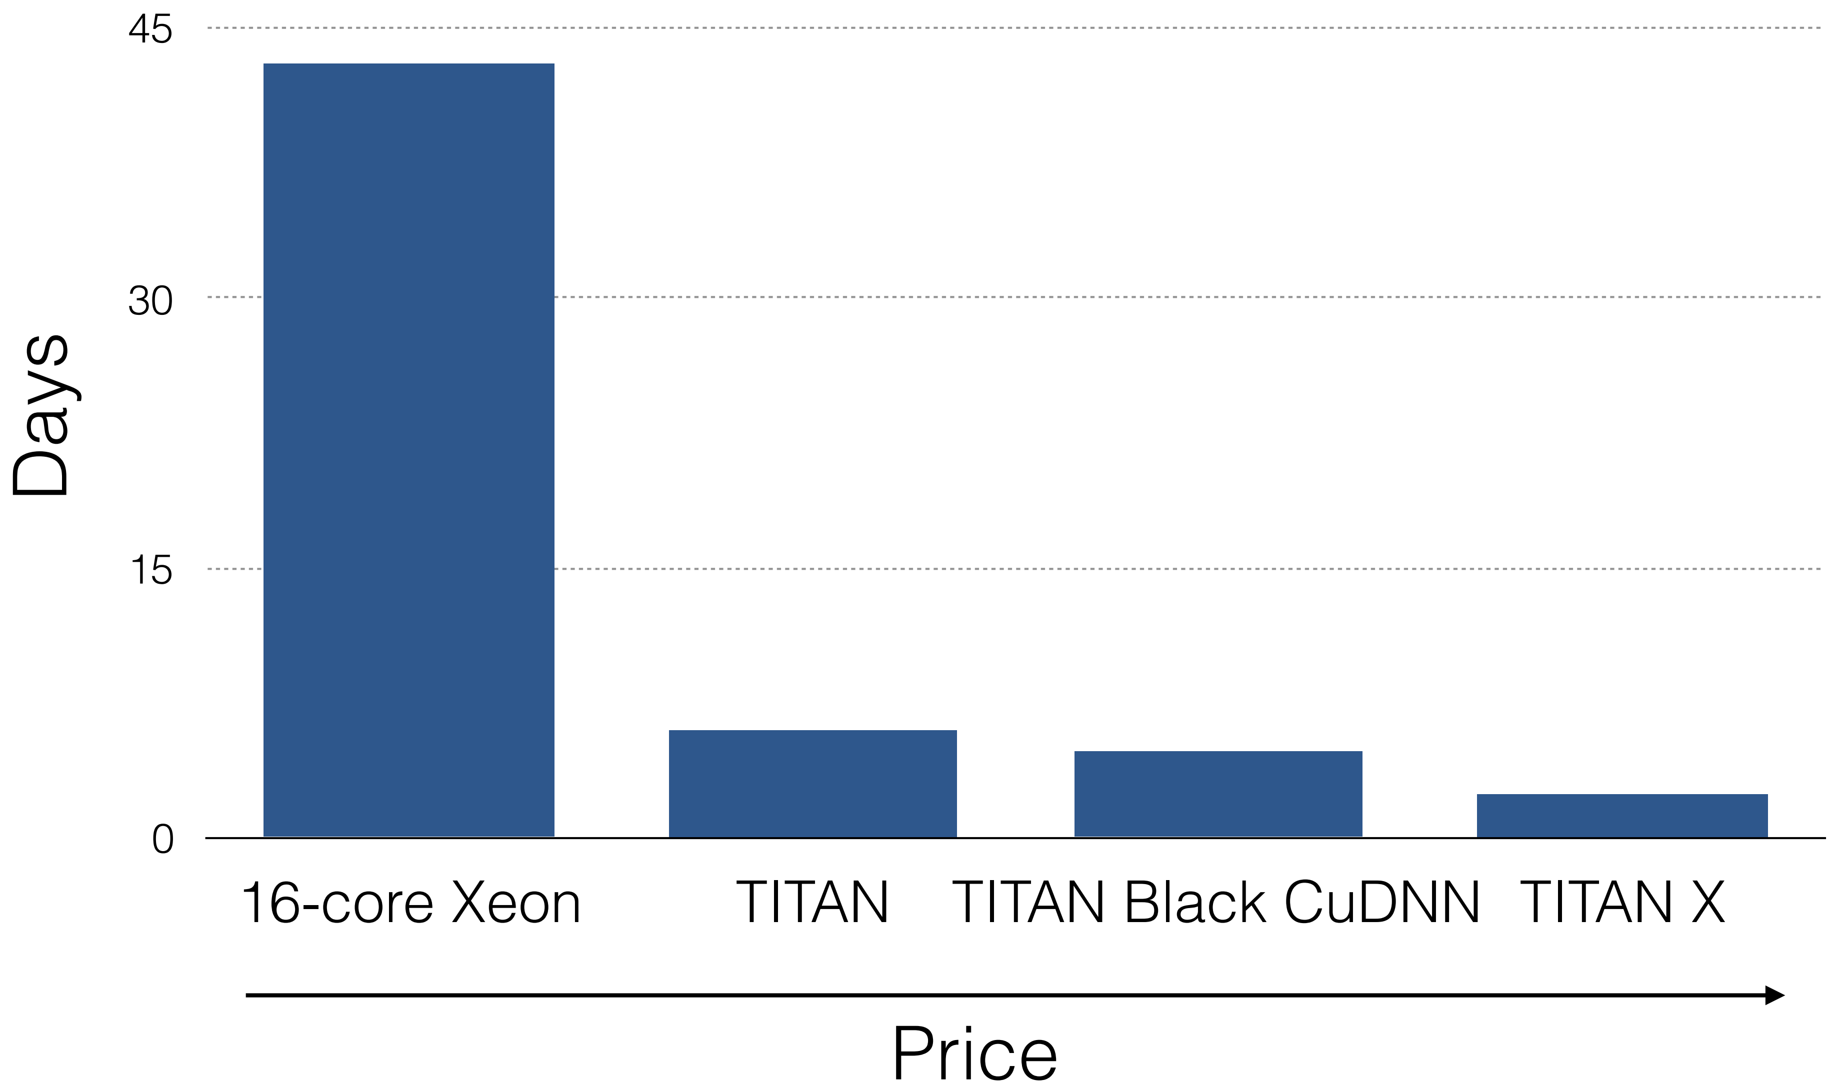

Source: NVIDIA

Rethinking the problem

- Make samples better with high-contrast tissue prep.
- Sparse reconstruction vs. dense reconstruction
- Avoid overfitting with unsupervised methods
- Cell-specific reconstruction strategies
- Solve this problem more like people do

Assisted Reconstruction Technique for Electron Microscopic Interrogation of Structure (ARTEMIS)

(a) tagging a genetically identified cell with an electron-dense tracer

(b) enhancing the electron-dense staining of these tracers

(c) imaging the cell rapidly at relatively low resolution

(d) re-imaging small volumes at higher resolution to map connectivity

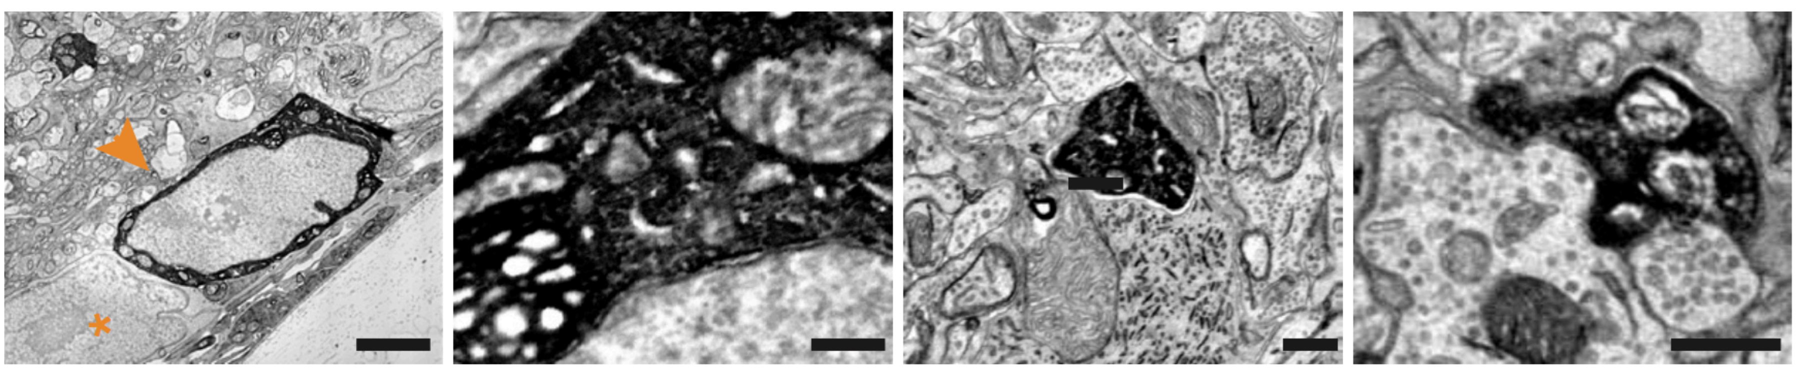

M. Joesch, D. Mankus, M. Yamagata, A. Shahbazi, R. Schalek, A. Peleg, M. Meister, J. W. Lichtman, W. J. Scheirer, J. R. Sanes, "Reconstruction of Genetically Identified Neurons Imaged by Serial-Section Electron Microscopy," eLife, Vol. 5, e15015 2016

ARTEMIS Staining Strategy

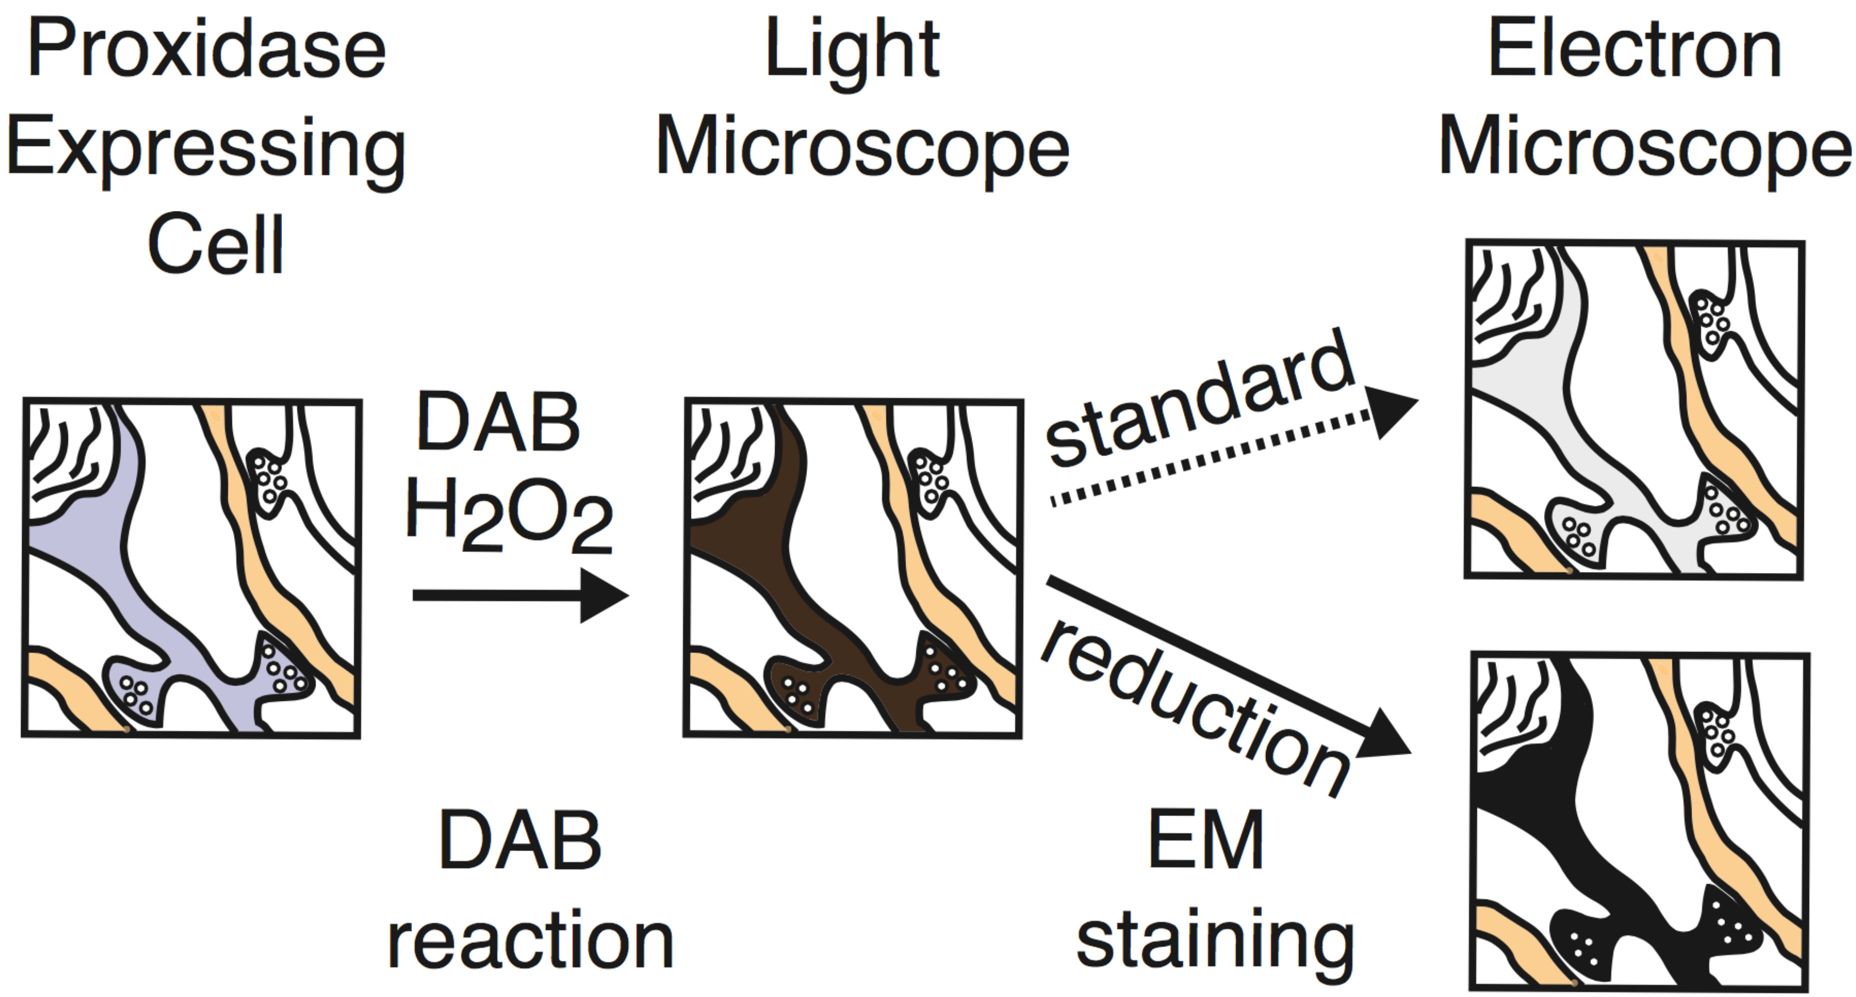

Improved contrast to noise ratio

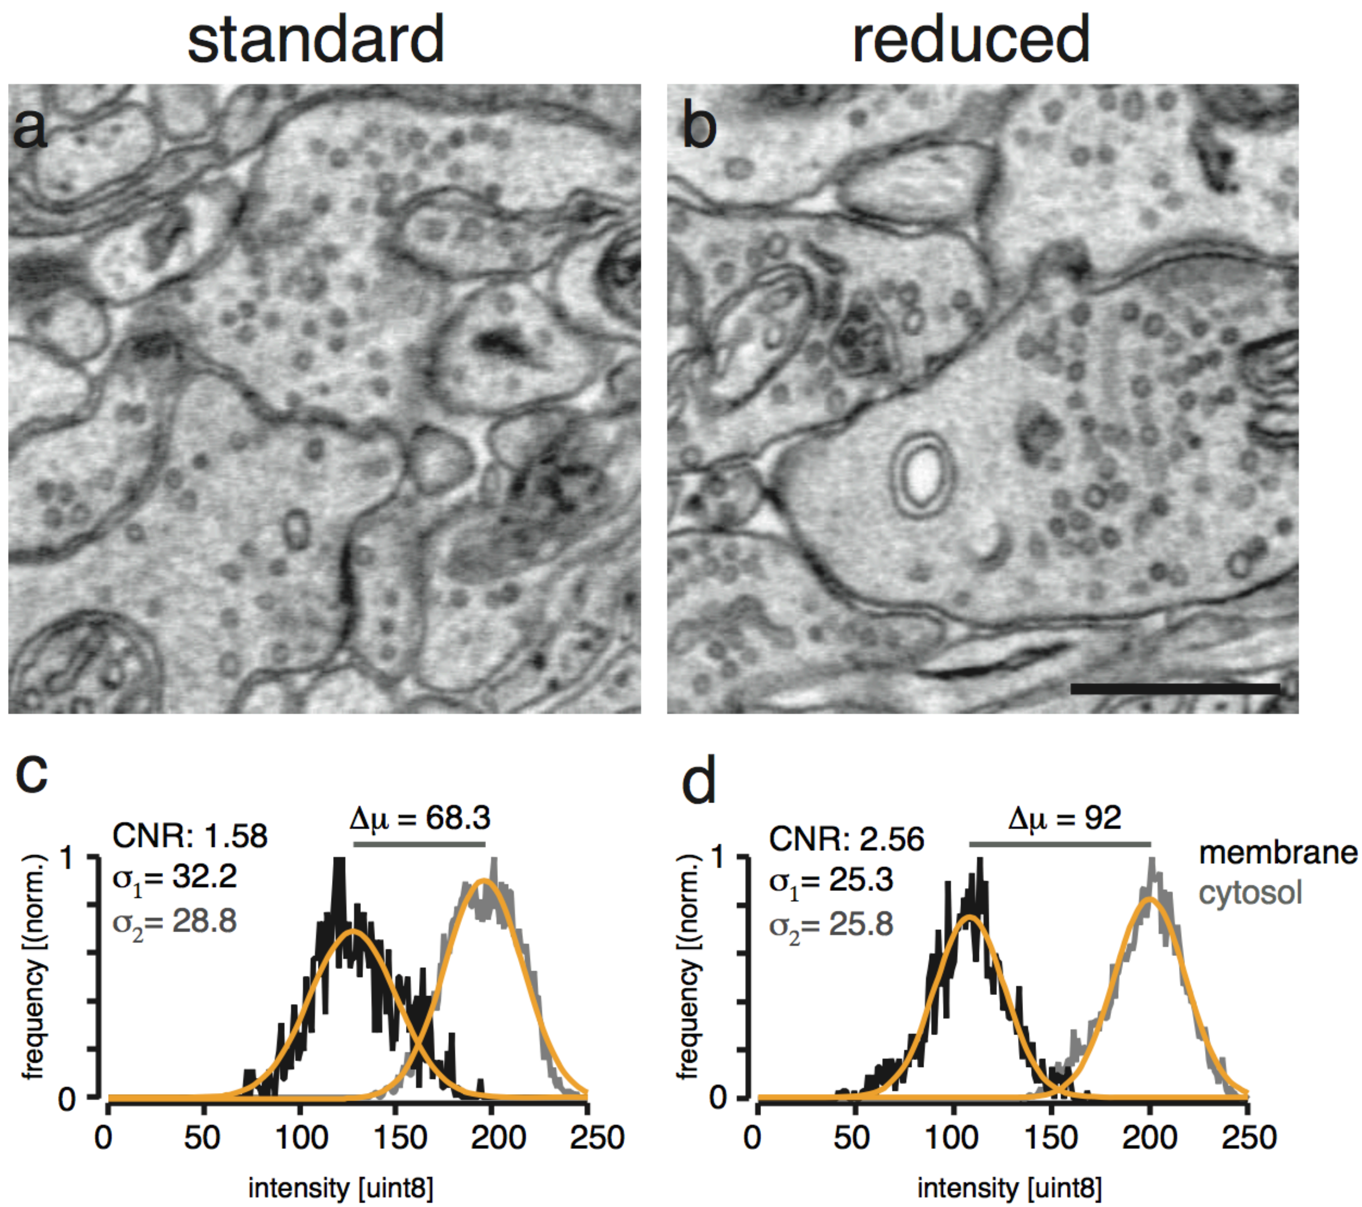

Cytosolic Apex (enhanced tissue)

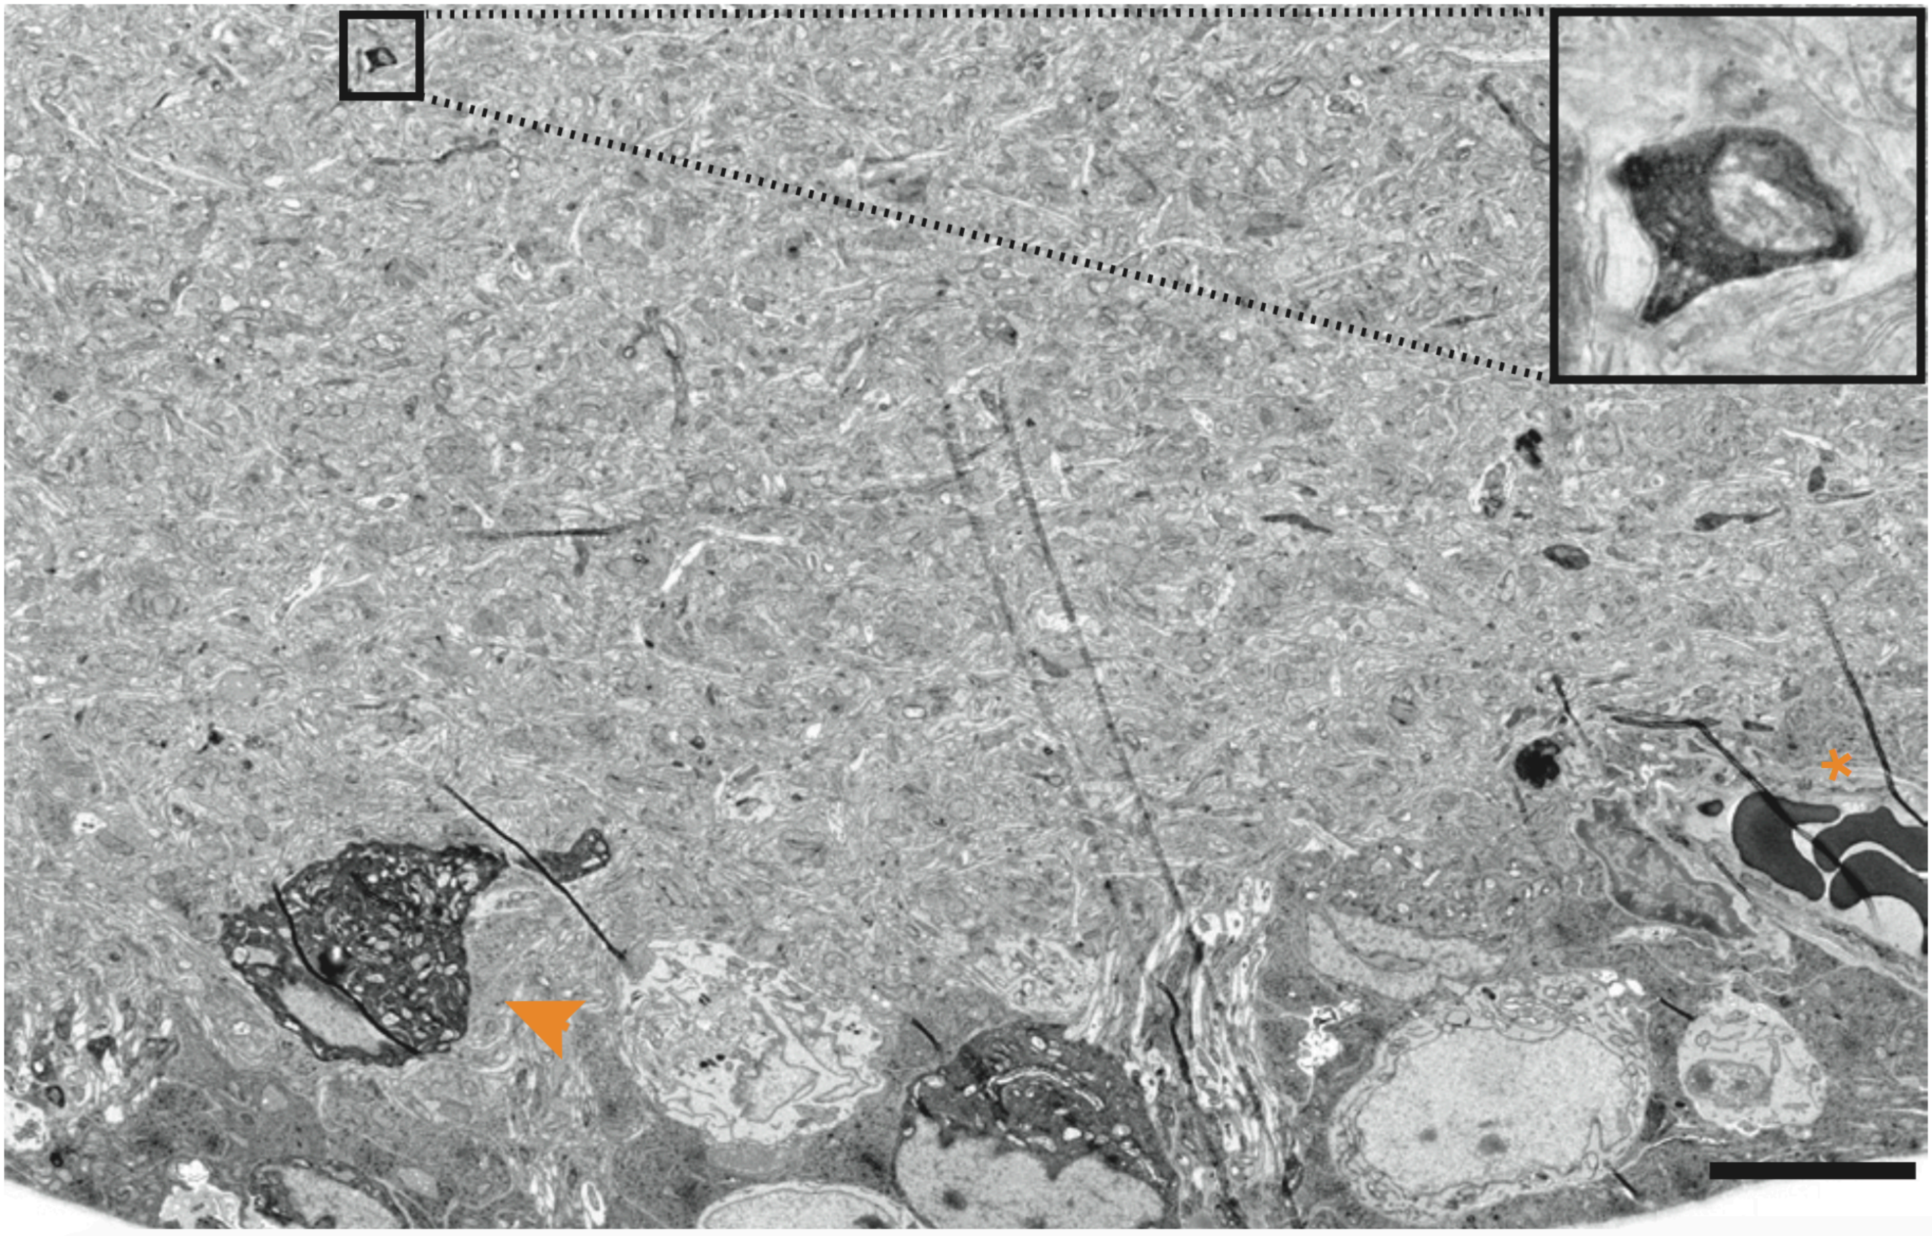

4nm/pixel, detail from 680µ x 680µ section

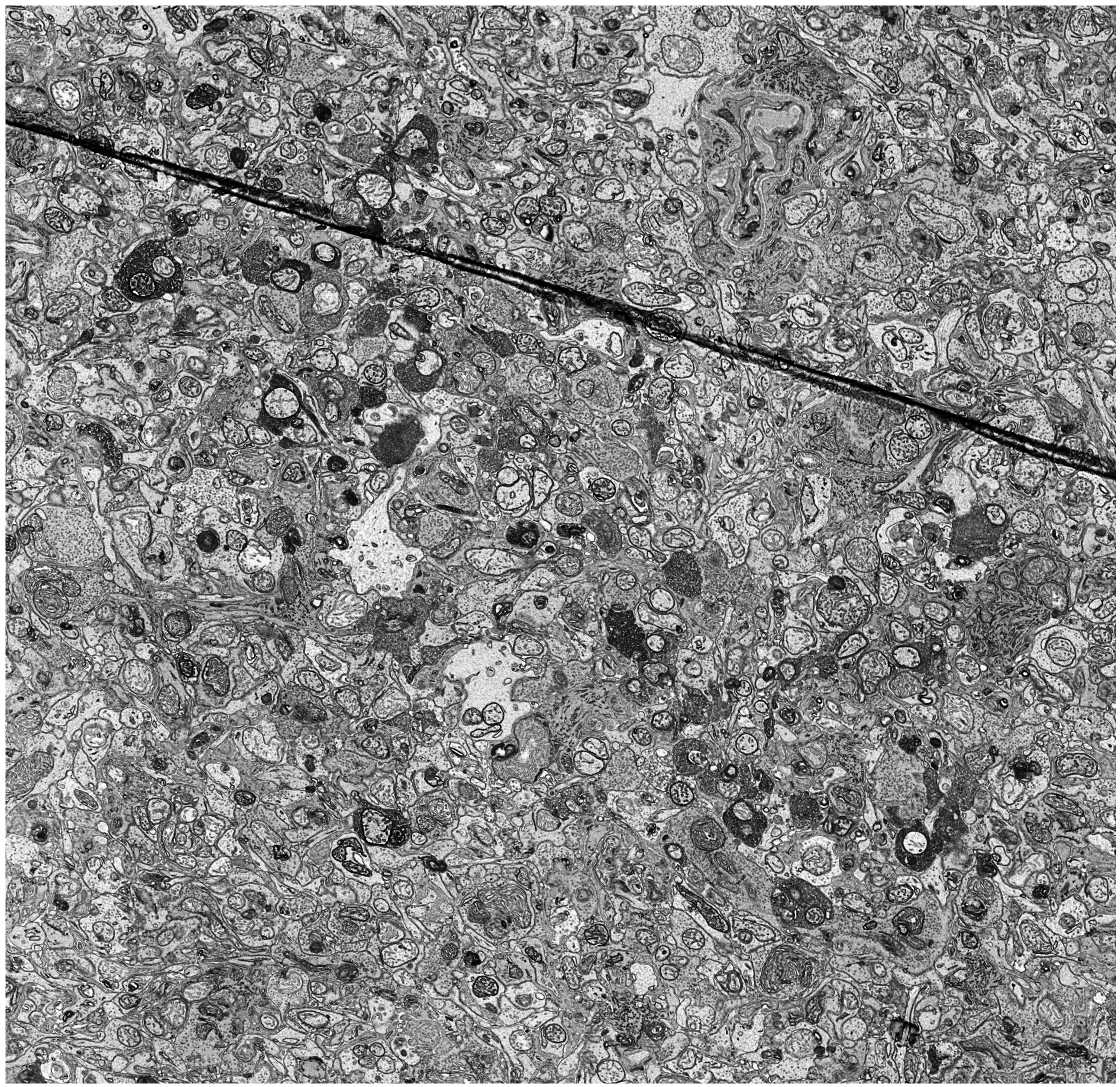

Learning-free 2D segmentation and 3D reconstruction

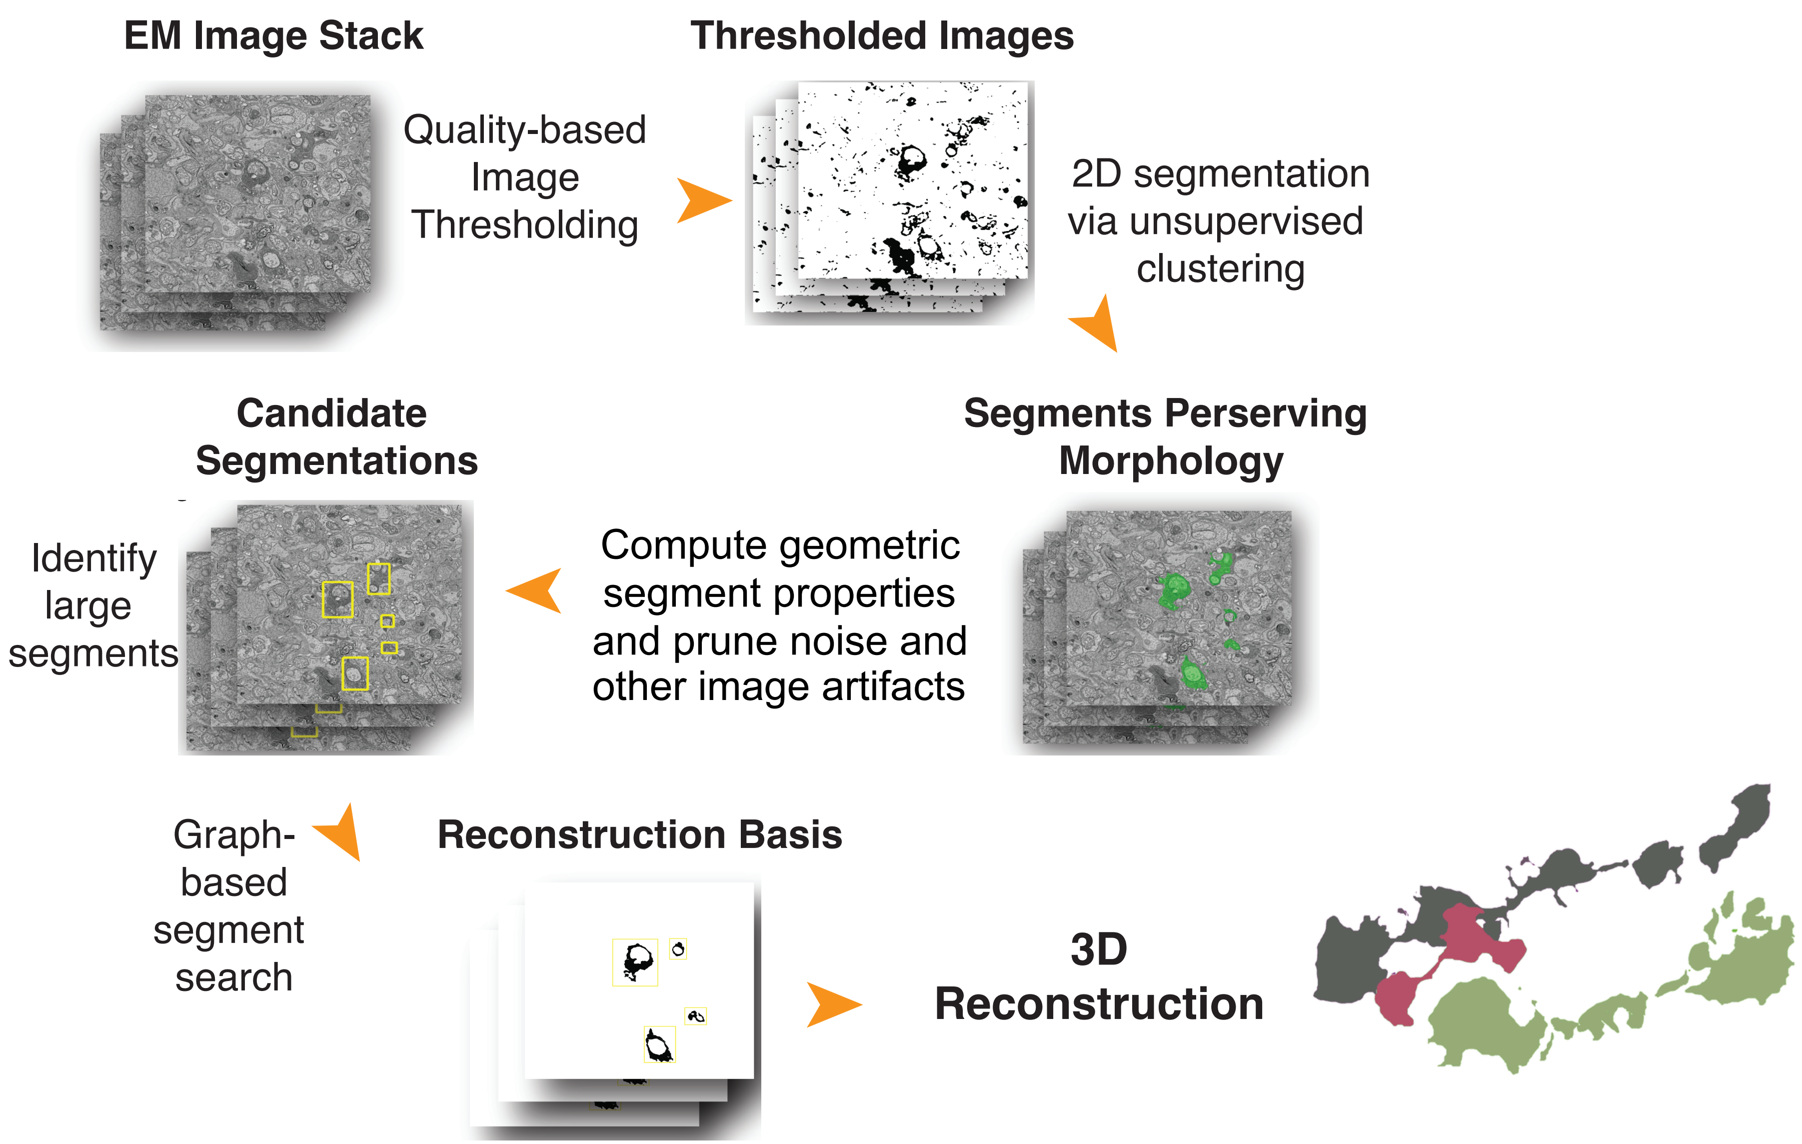

Starburst Amacrine Cell

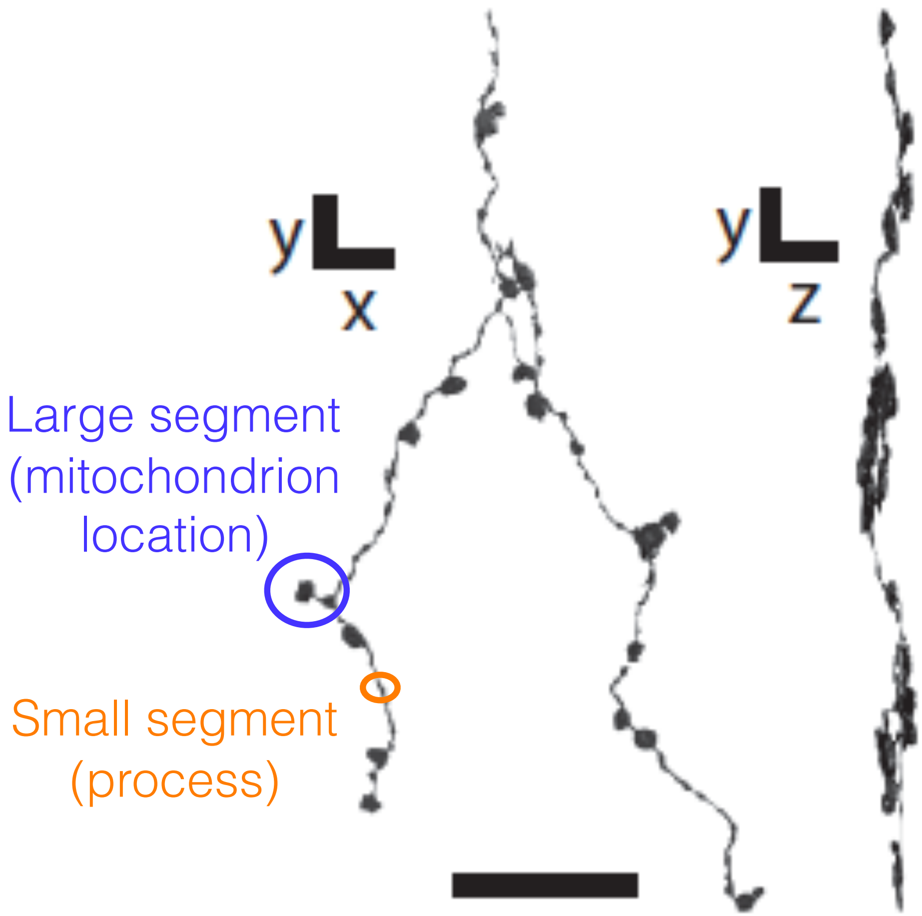

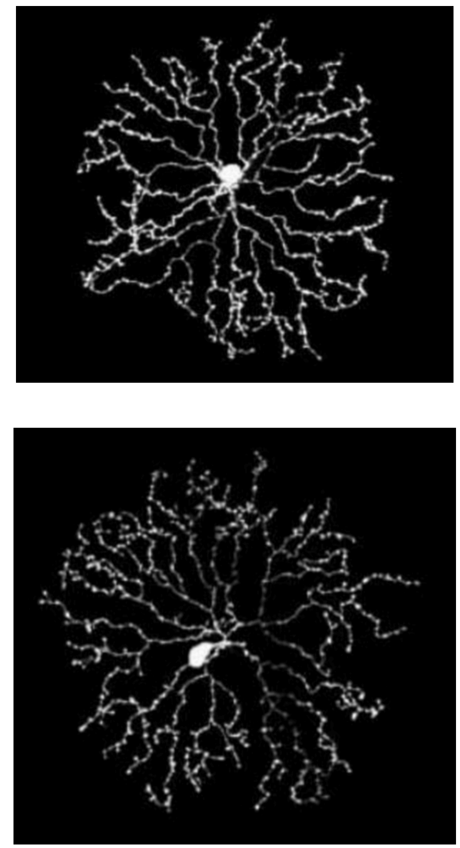

Keeley et al. J. Comp. Neurol. 2007

Step 1: Pre-processing

Original EM images are of very high resolution (100000 x 50000 pixels)

Process local neighborhoods in 2048 x 2048 tiles

Each tile can be processed separately for 2D segmentation

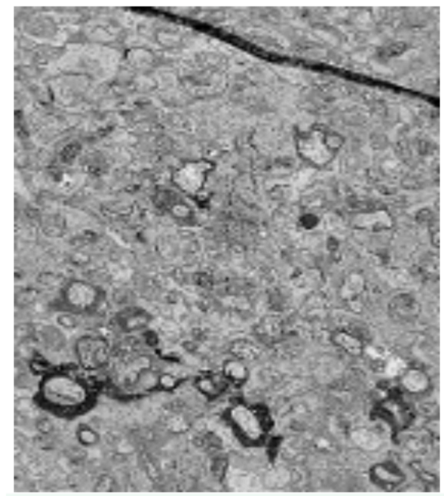

Step 2: Adaptive Thresholding on Local Intensity

Key step enabled by ARTEMIS

Apply a Wiener filter, modulating via known pixel ranges of ARTEMIS markers:

noise variance

pixel

$$\beta(n_1, n_2) = \mu + \frac{\sigma^2 - \nu^2}{\sigma^2} \times (a(n_1, n_2) - \mu)$$

pixel mean

pixel variance

Step 3: Cluster-based Image Segmentation

Choose threshold to minimize the intraclass variance of the black and white pixels

Prune non-ARTEMIS marked segments: 3 x n matrix to calculate weights based on local pixel neighborhoods

Remove Artifacts

Step 4: Assess morphology

Catalog the properties of each segment:

- Centroid
- Convex Area
- Area
- Bounding Box
- Extent
- Solidity
- Extrema
- Major Axis Length
- Minor Axis Length
- Equiv. Diameter
- Eccentricity

Step 5: Database assessment

Analysis moves to abstract representations stored in a MySQL database

Rapid search for large and small segments repeated across layers

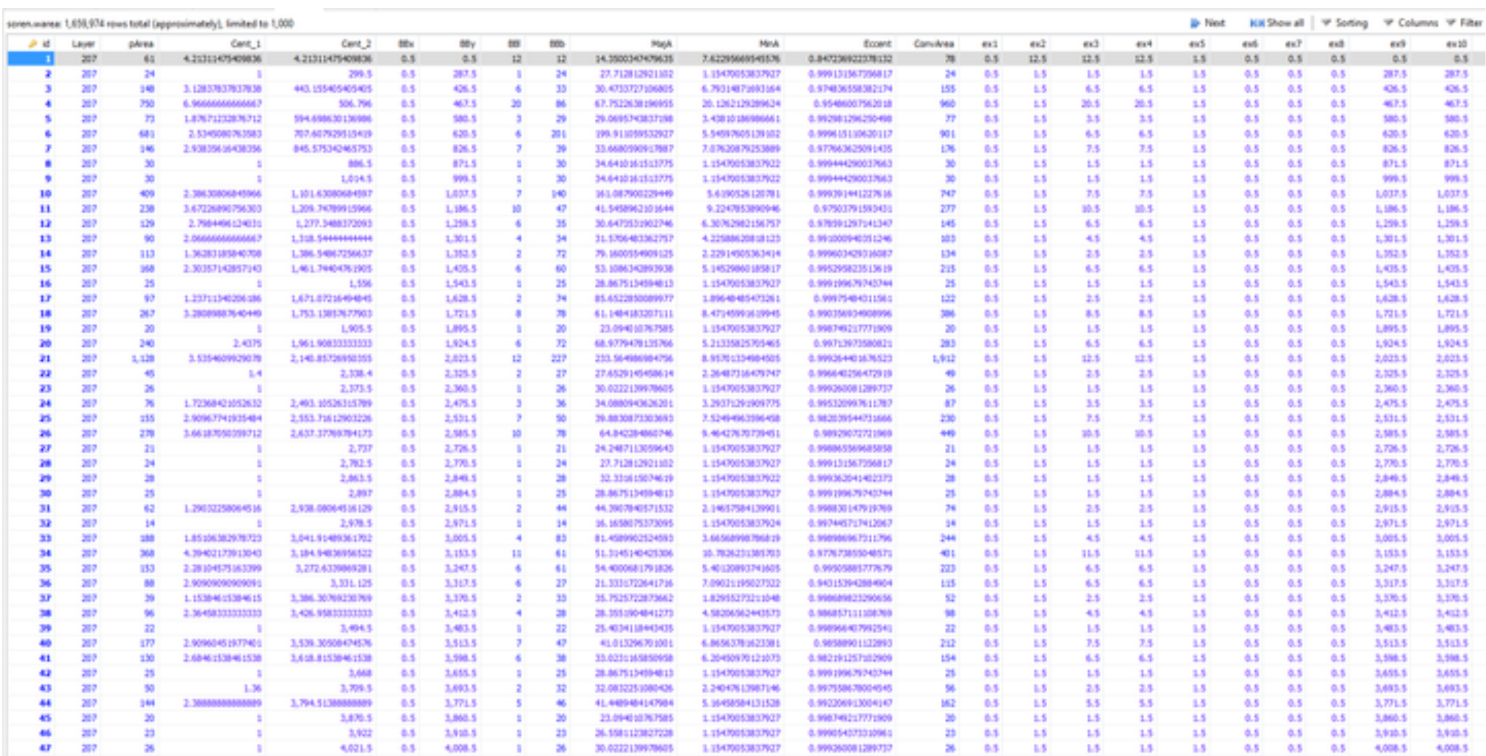

Step 6: Graph-based segment search

Greedy strategy:

Starting from larger segment, exhaustively reconstruct within a bounding box for a limited number of layers

Register all layers into one 2D candidate

Identify direction of the process, and then expand bounding box according to it.

Repeat steps until the connectivity path reaches another large segment or another smaller path

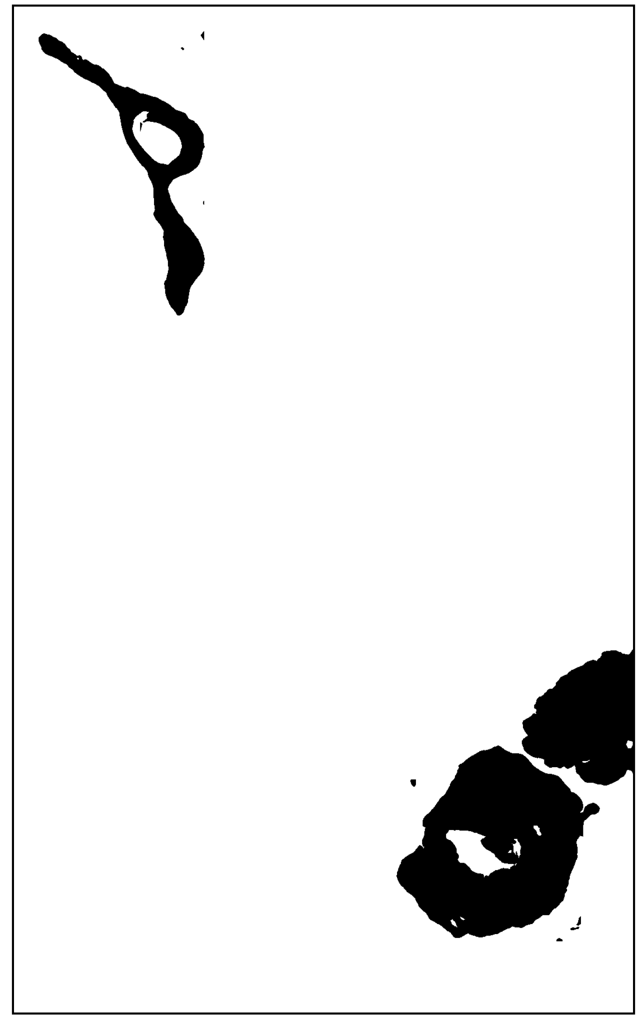

Reconstruction Results

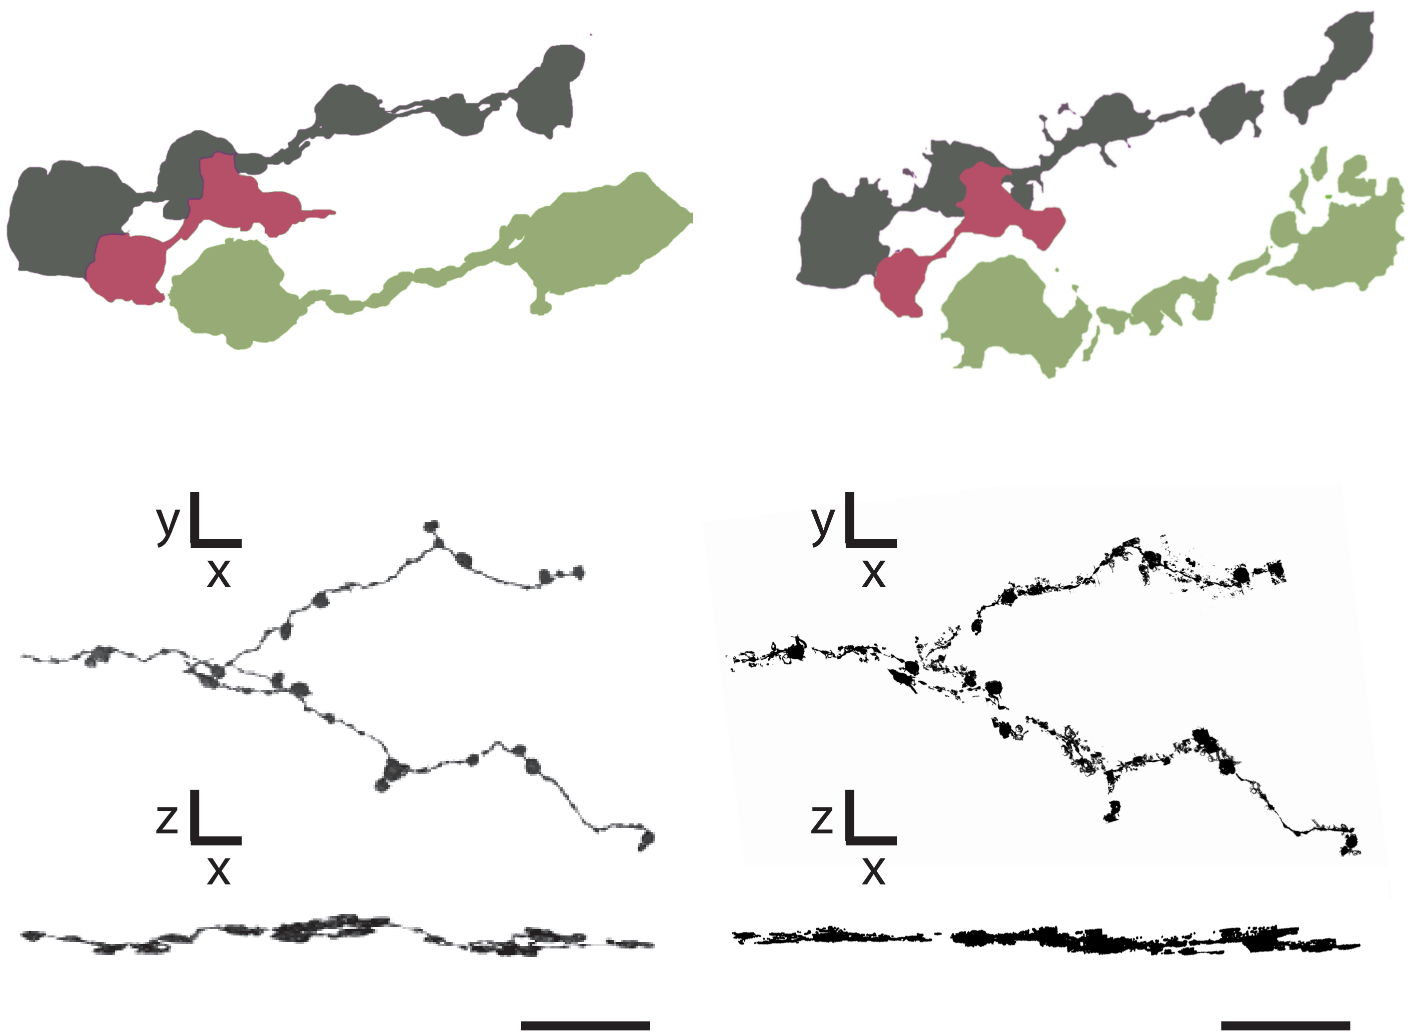

Reconstruction Results

Ground Truth | Our Method | Ground Truth | Our Method

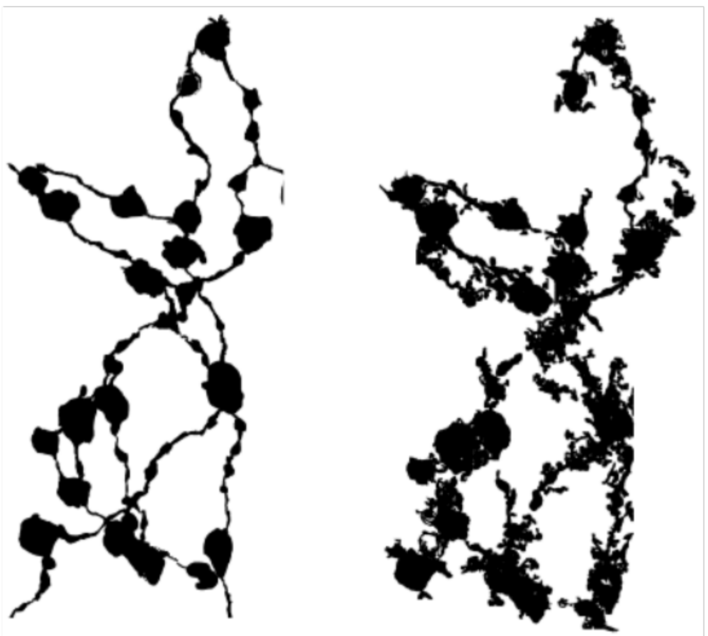

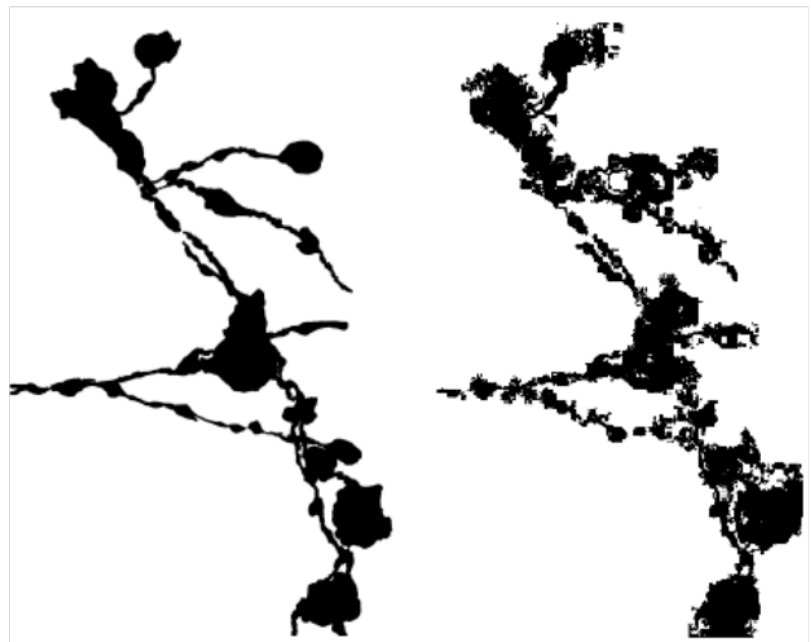

Quantitative Performance

Ground-truth: 444 APEX positive segments

Recall statistic: 91.8% for the 2D segmentation portion of the algorithm

Two orders of magnitude faster compared to Random Forests* (supervised machine learning)

Not perfectly accurate, but **much** faster

*Kaynig et al. Medical Image Analysis 2015

Quantitative Performance

	3D U-Net	Ours (Cluster Based Thresh.)	Ours (Adaptive Thresh.)
Manual Annotation	40 Hr.	-	-
Training	926 Min.	-	-
Segmentation	233 Min.	8 Min.	8 Min.
Reconstruction	-	128.8 Min.	16.82 Min.

Also works for X-ray

Original

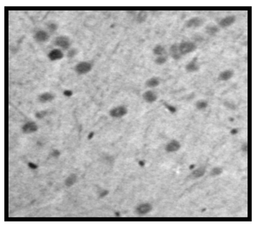

Contrast Enhance

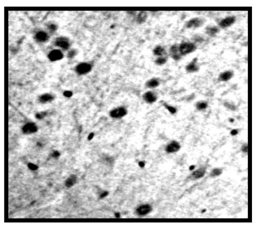

Adaptive Thresh.

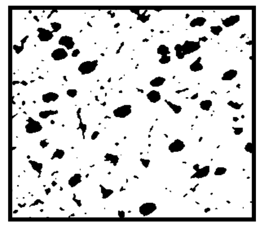

Segment Enhance.

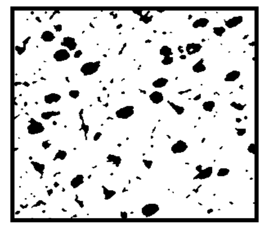

Artifact Removal

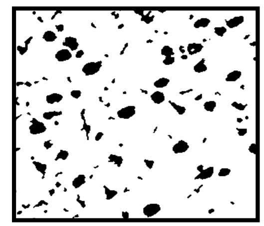

Learning-free Segment.

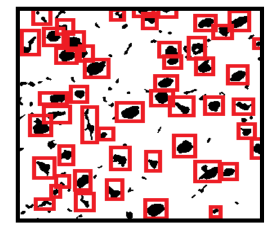

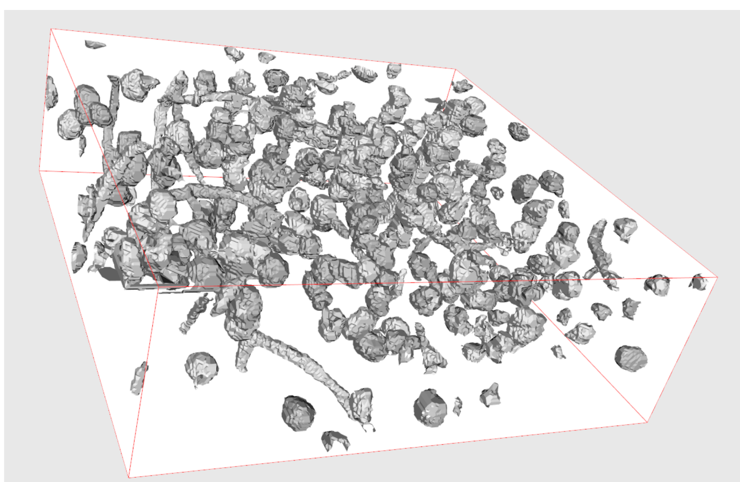

Reconstruction

Reconstruction Results

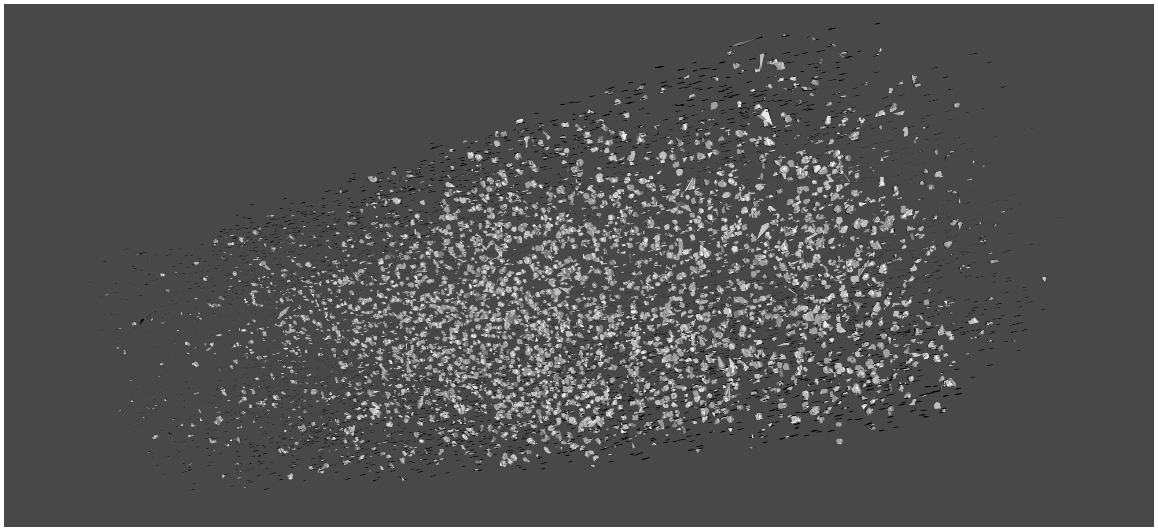

Cells

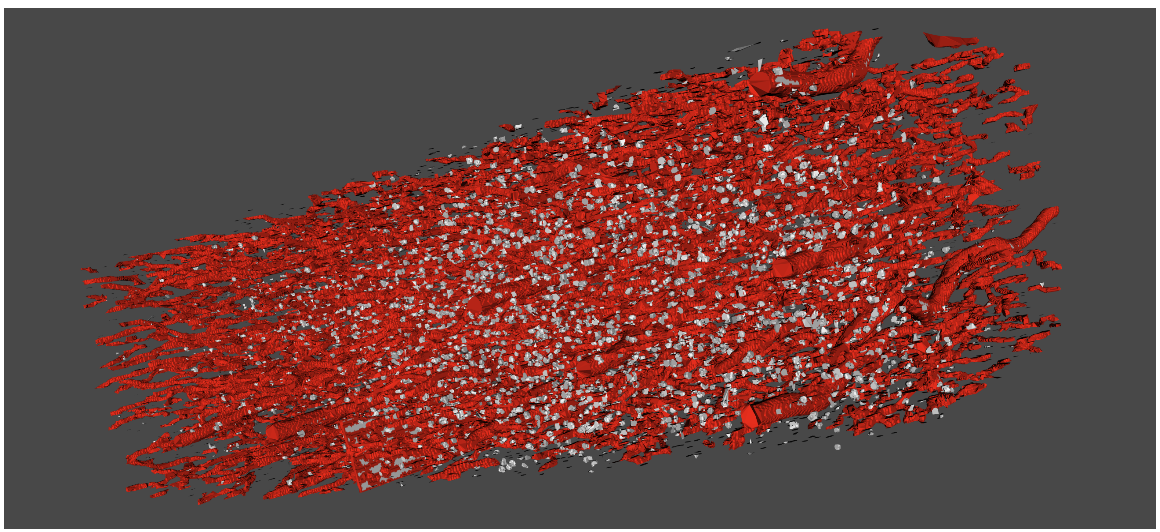

Vasculature + Cells

Apex Cells Imaged with X-ray

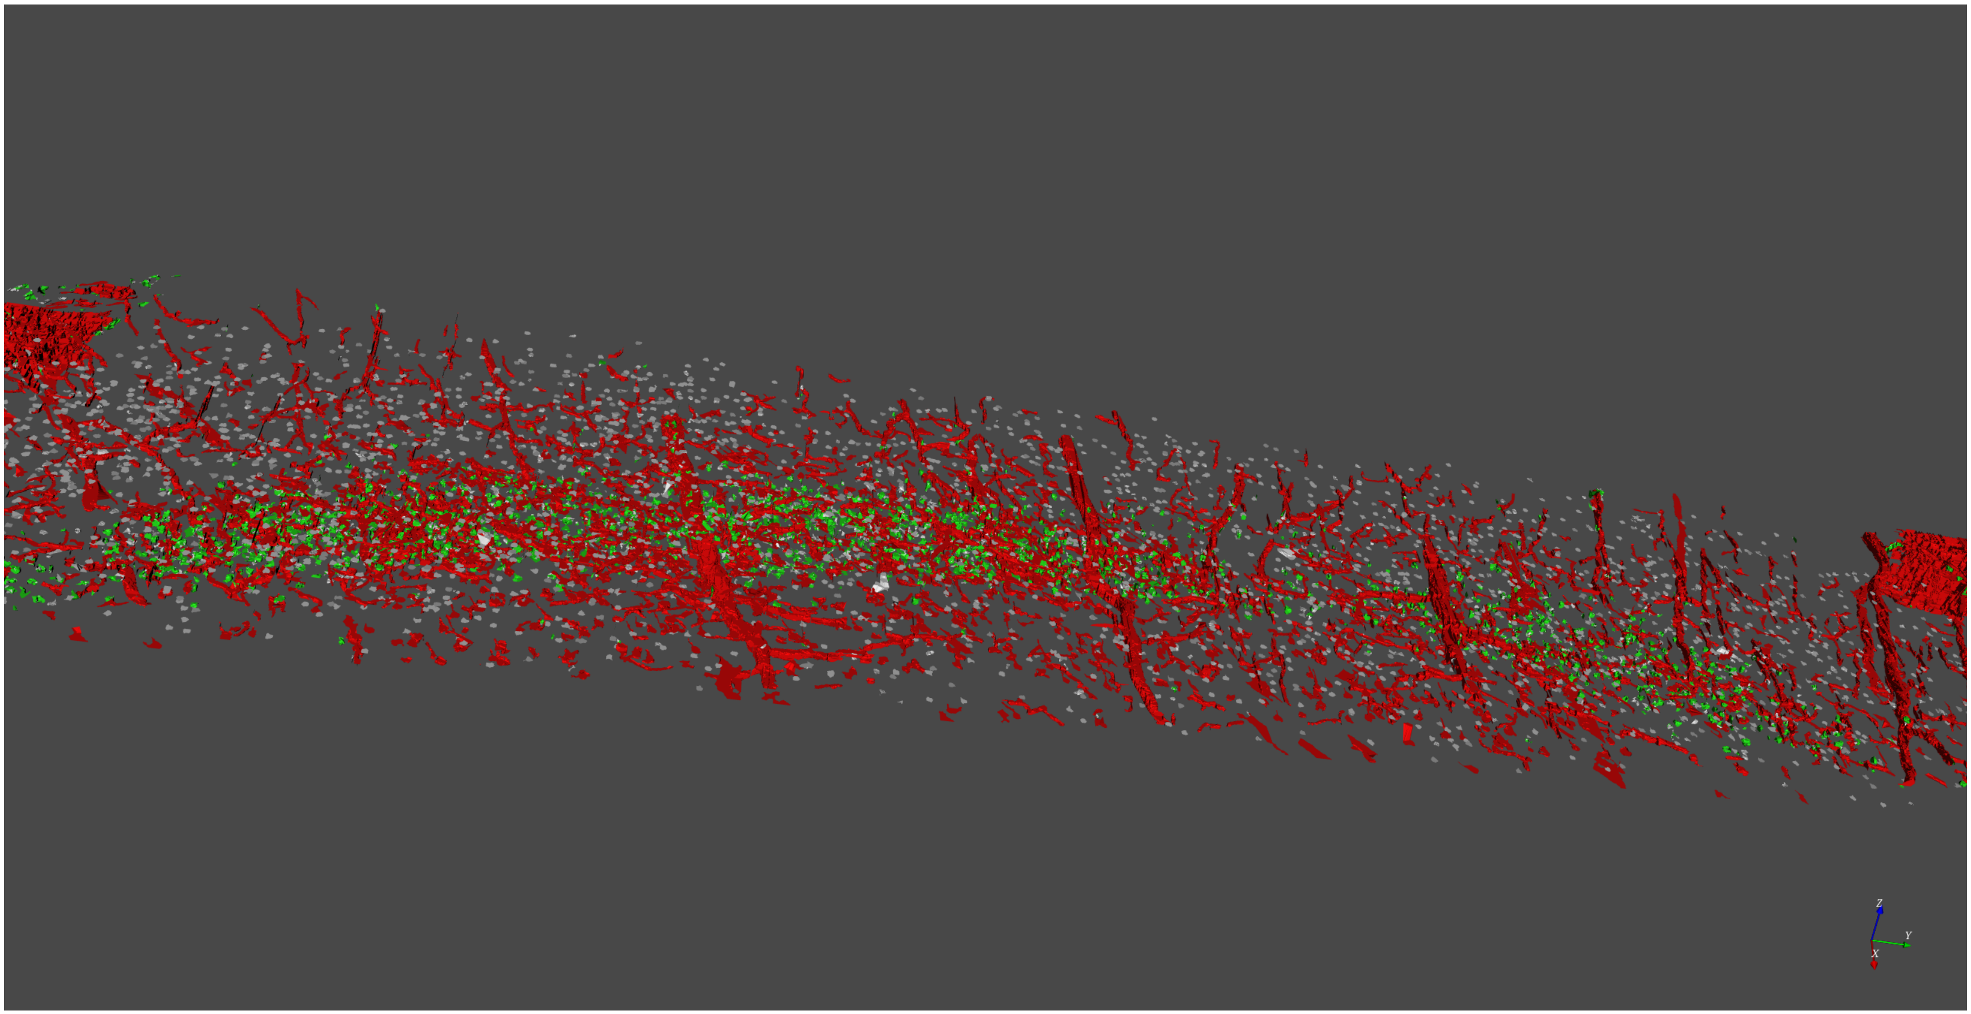

Comparison to deep learning

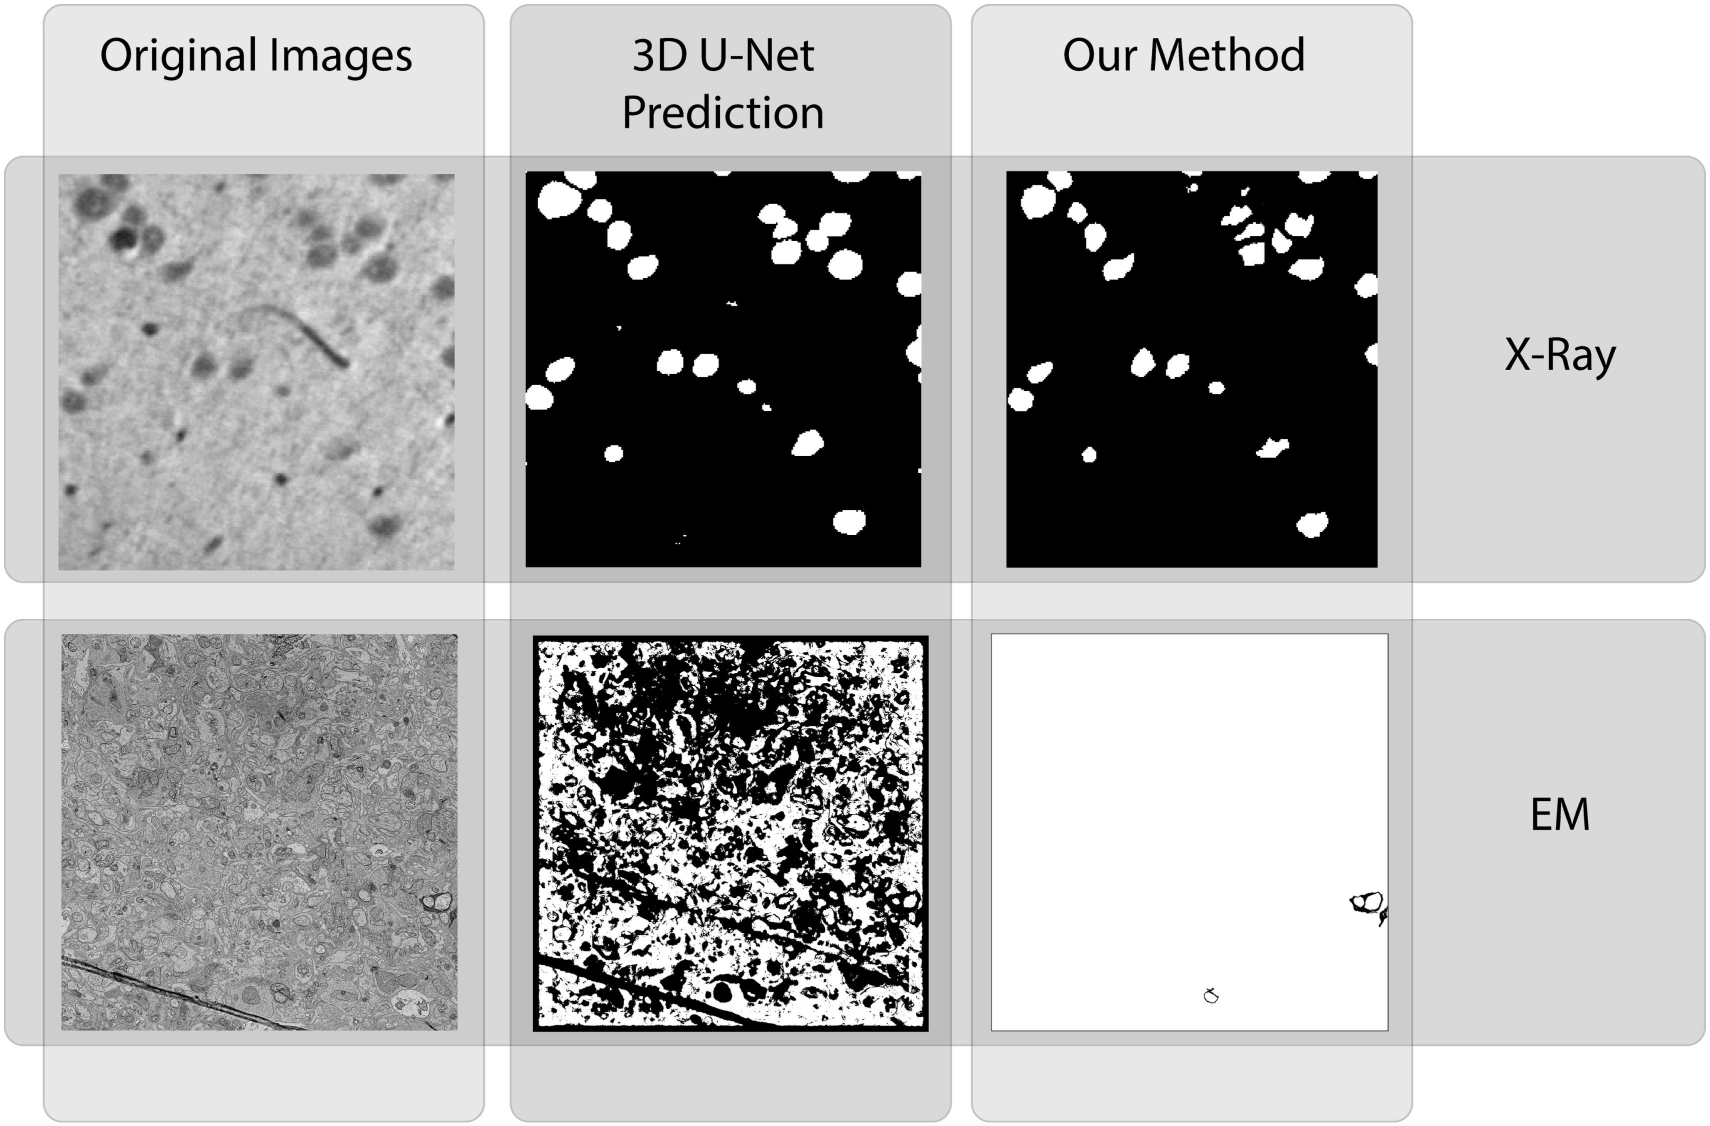

Bringing learning back in

Open Set Machine Learning

Can we bring supervised machine learning back into the picture to handle unknown data?

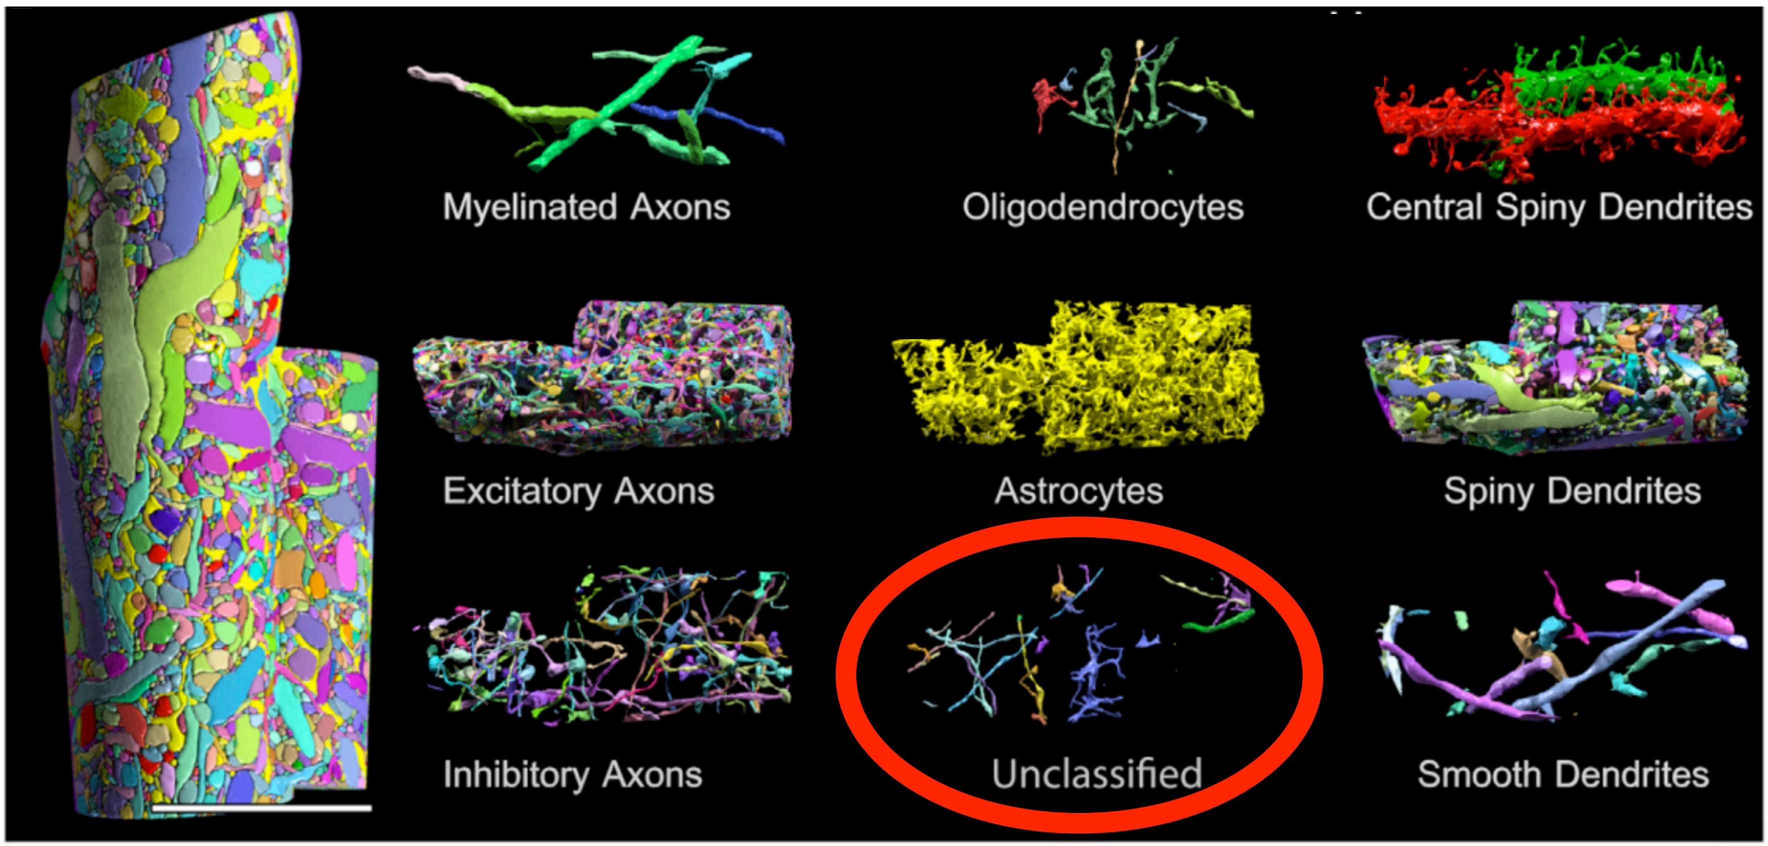

Kasthuri et al. Cell 2015

“There are known knowns…”

known classes: the classes with distinctly labeled positive training examples (also serving as negative examples for other known classes)

known unknown classes: labeled negative examples, not necessarily grouped into meaningful categories

unknown unknown classes: classes unseen in training

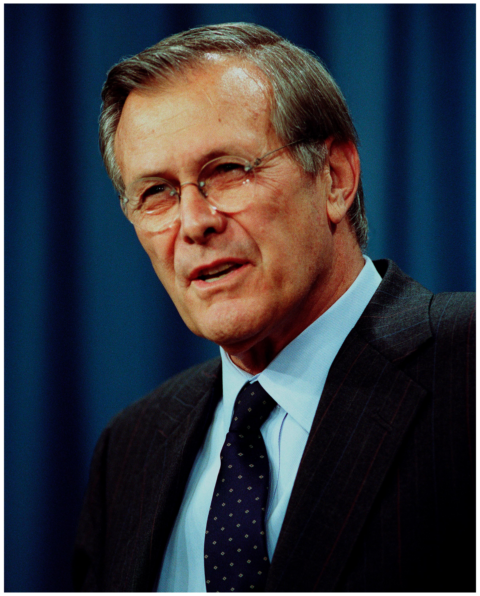

Learning Objective

Minimize open set risk:

Training Data

$$\underset{f \in \mathcal{H}}{\operatorname{argmin}} \left\{ R_{\mathcal{O}}(f) + \lambda_r R_{\mathcal{E}}(f(\hat{V} \cup \hat{K})) \right\}$$

Open Space Risk Associated with **U**

Regularization Constant

Empirical Risk Function

W. Scheirer et al. IEEE T-PAMI 2013

MICrONS use case: “synaptomics”

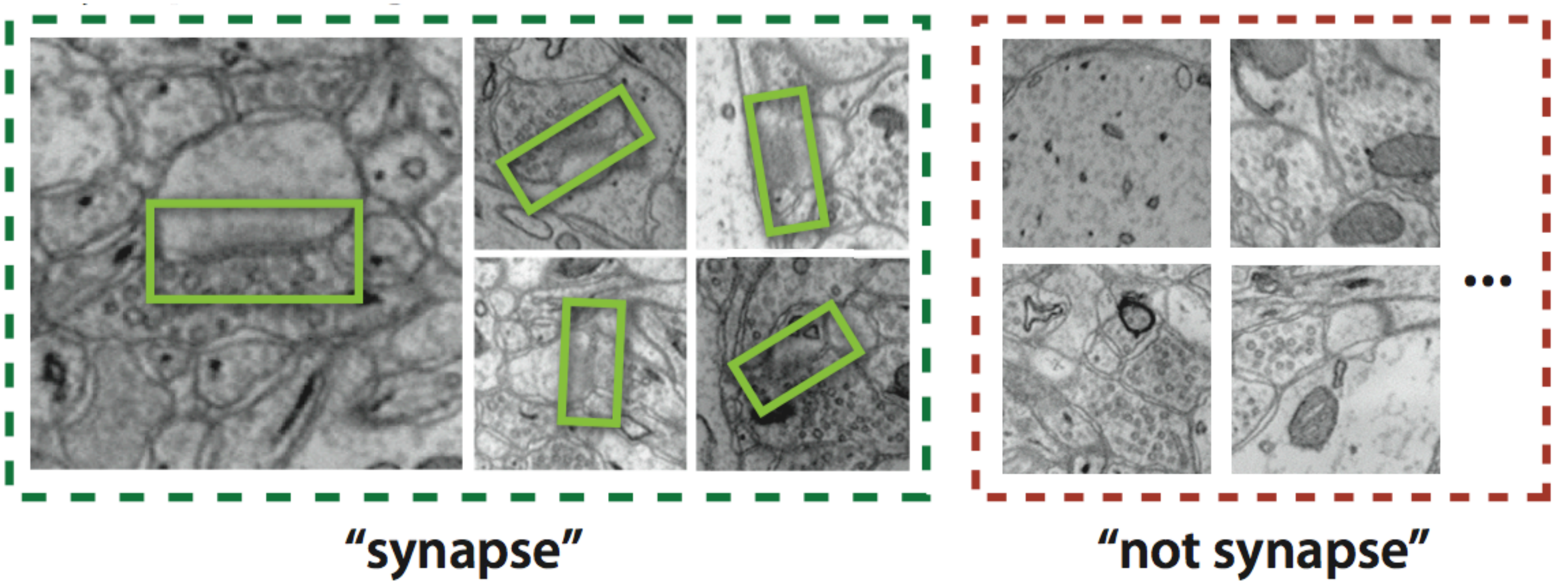

Model: Compact Abating Probability

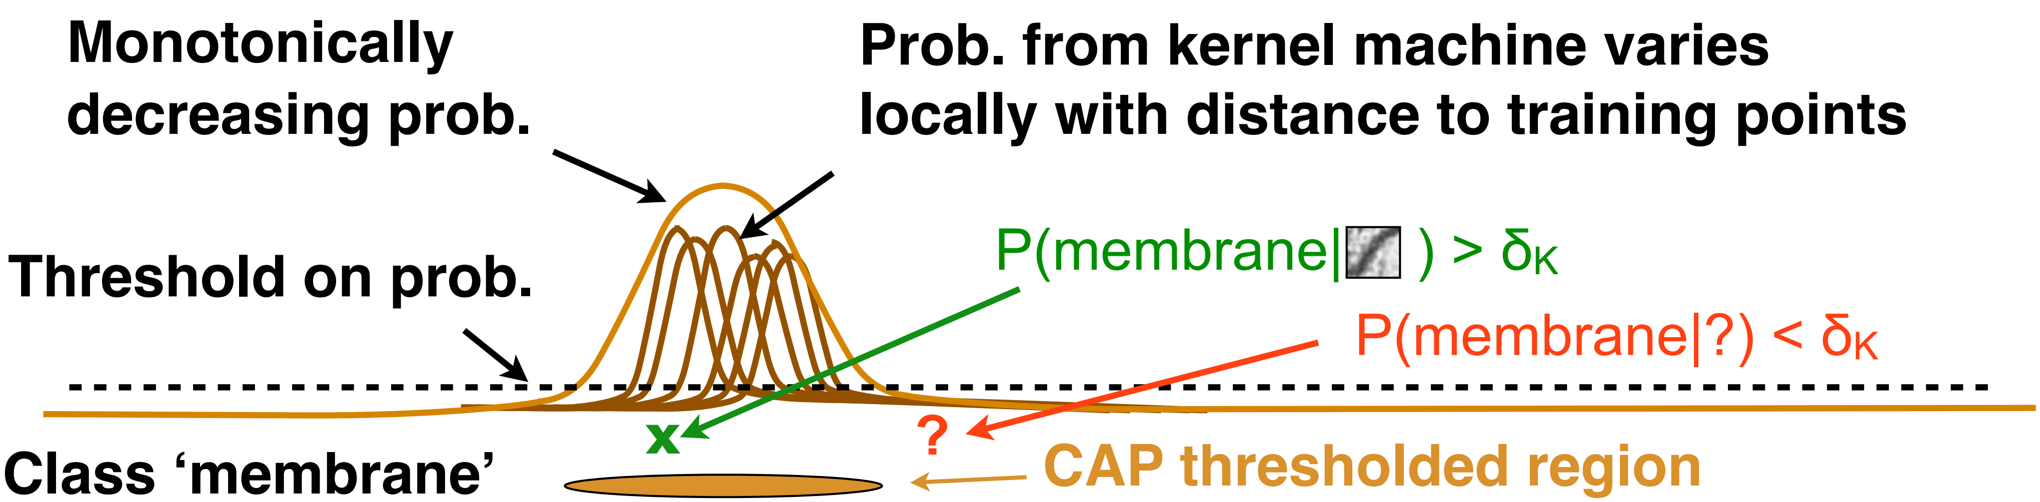

W. Scheirer et al. IEEE T-PAMI 2014

Binary RBF SVM incorporating a CAP model: W-SVM

Combine probabilities computed for both 1-class and binary RBF SVMs

1-class SVM CAP model is a conditioner

could be very small

if $P_O(y|x) > \delta_\tau$, then

consider $P_O(y|x)$

else

reject

Step 1: Train a 1-class SVM f^o

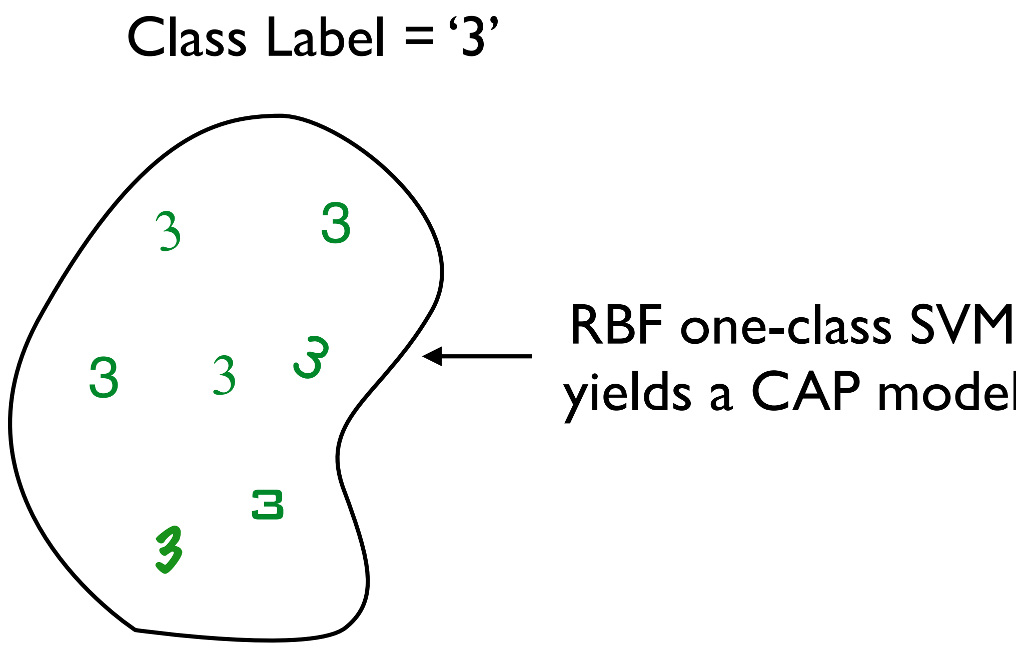

Step 2: Fit Weibull over tail of scores from f^{o}

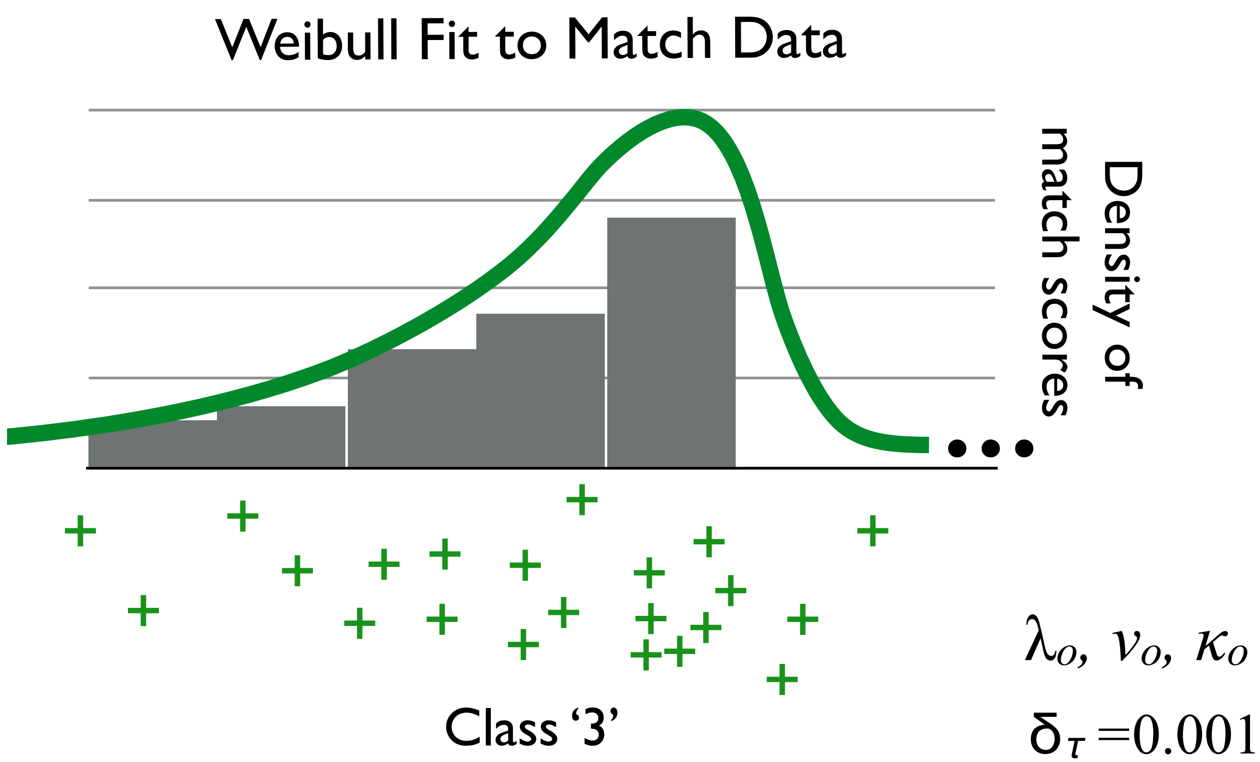

Step 3: Train a binary SVM f

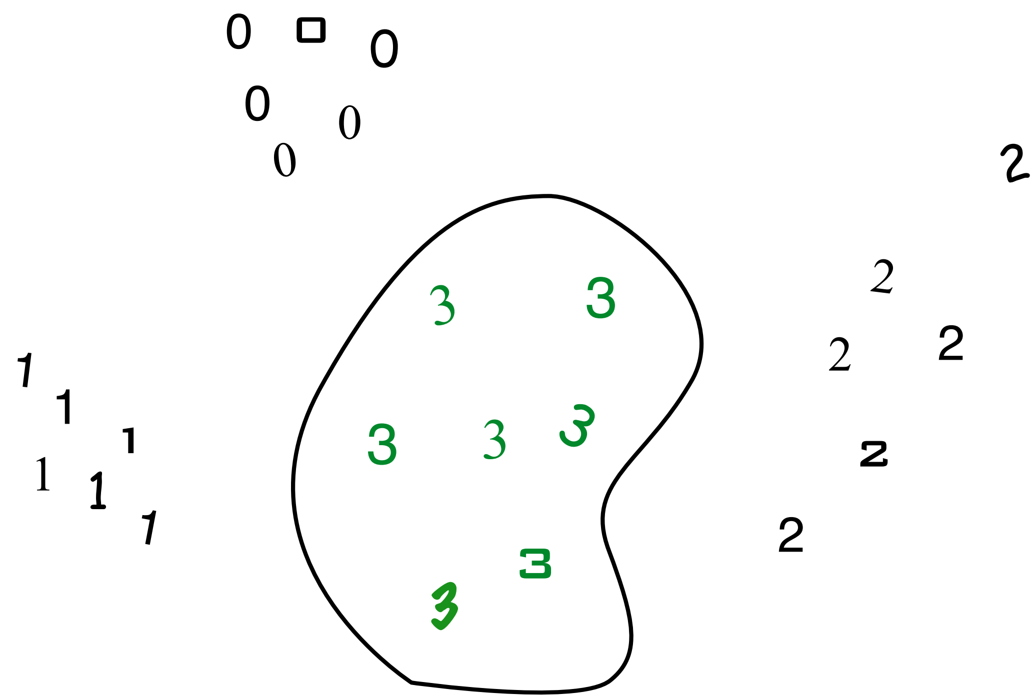

Class Label = '3'

Known Negative Classes = '0', '1', '2'

Step 4: Fit EVT distributions over tails of scores from f

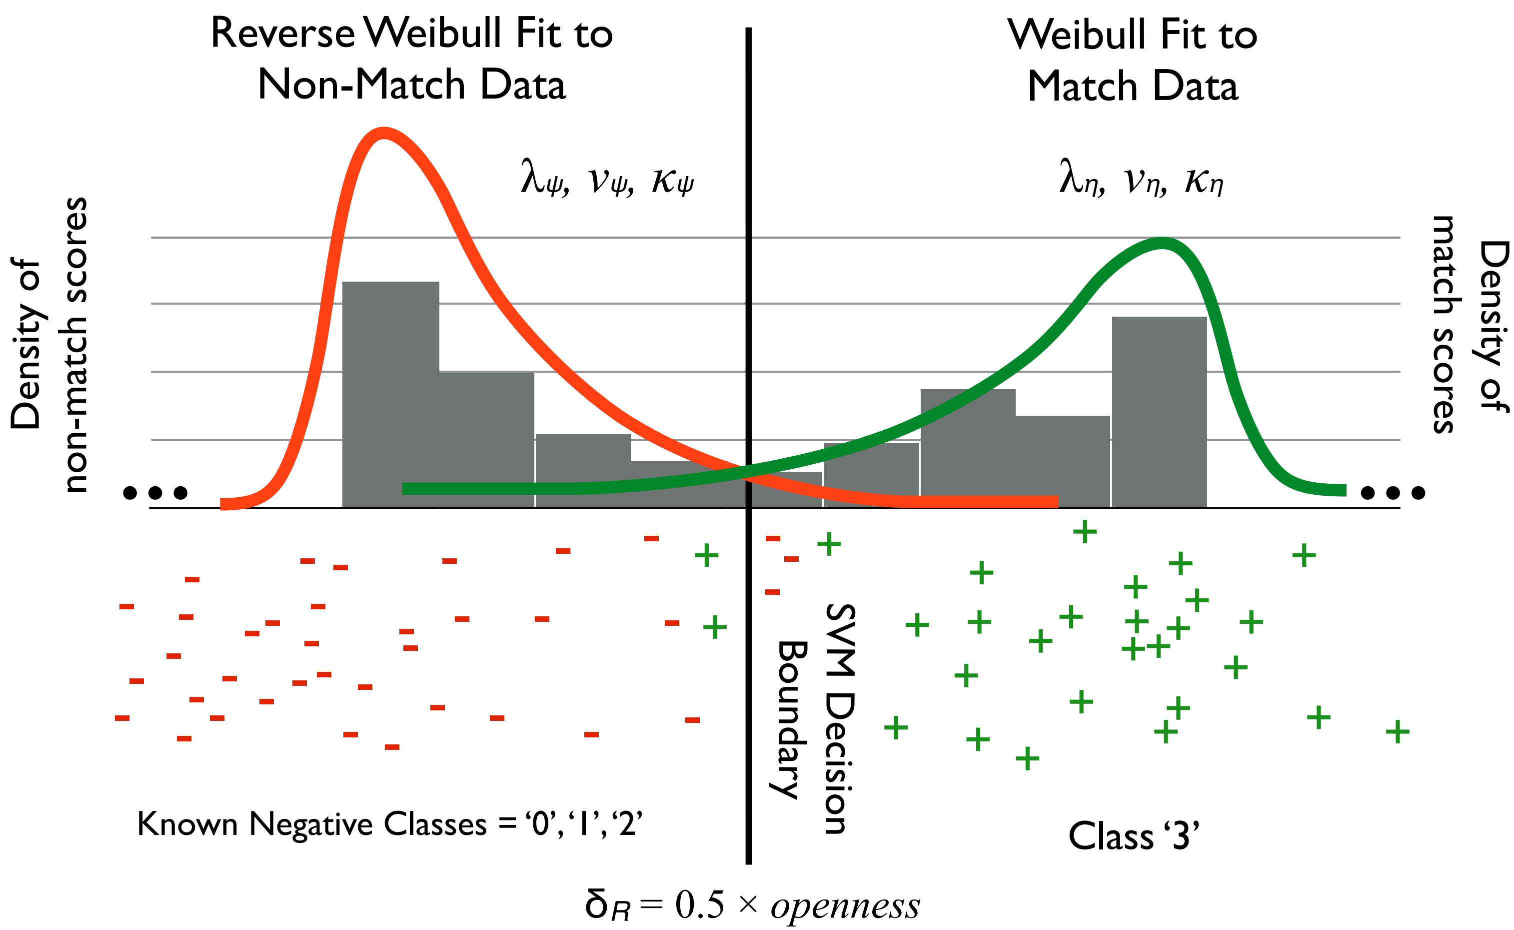

W-SVM Object Recognition

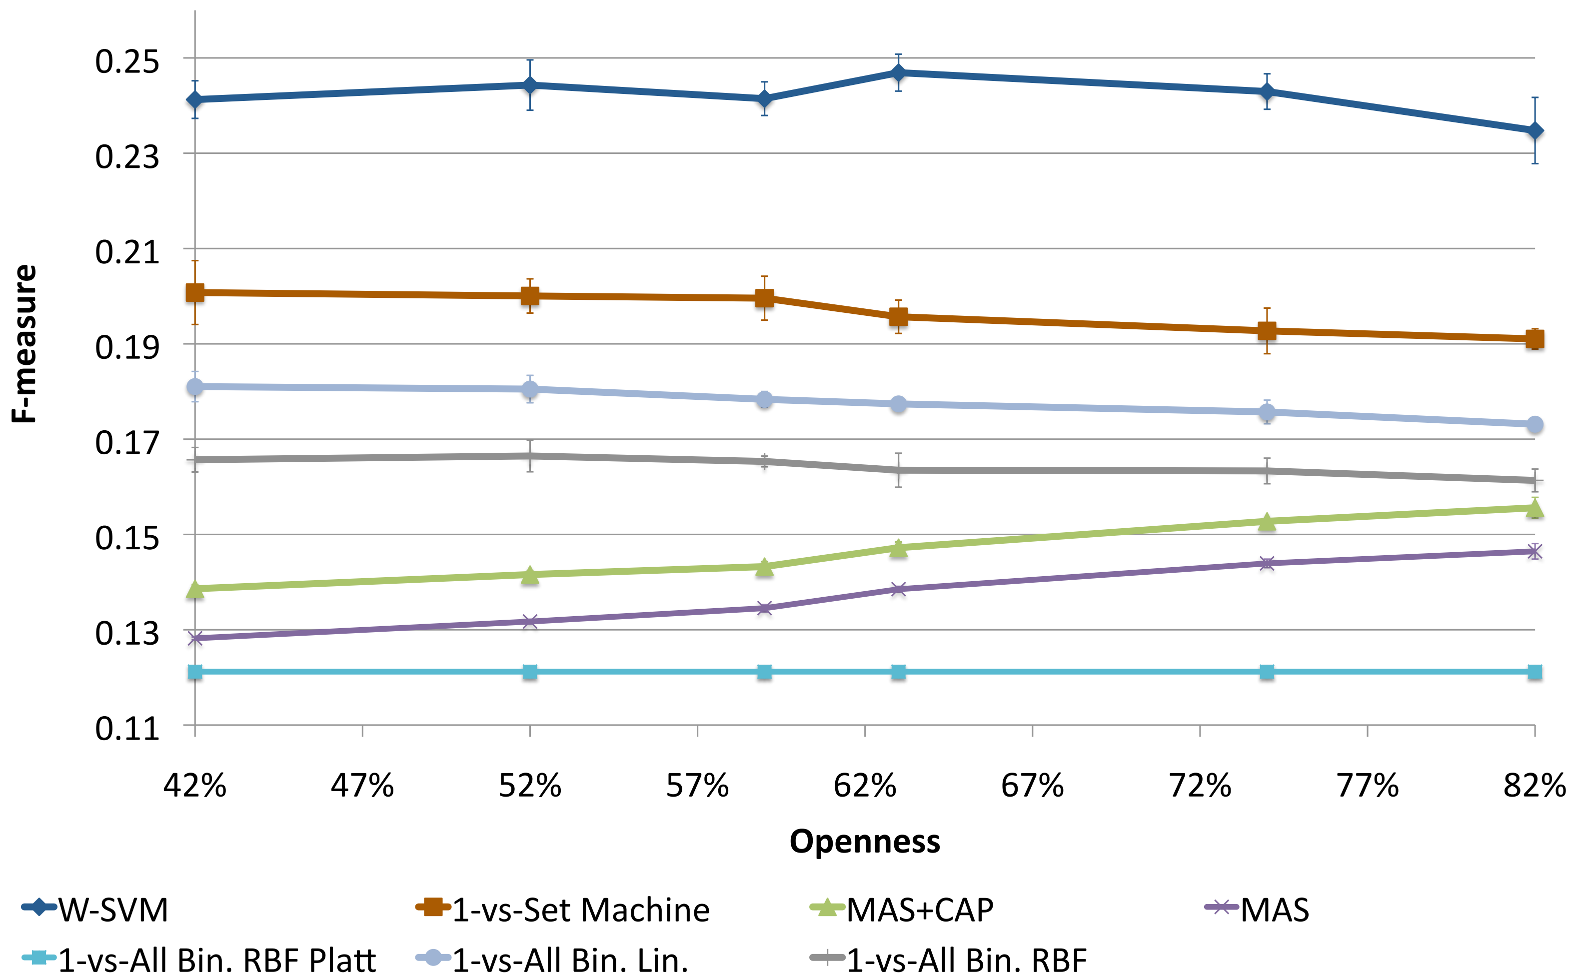

Incremental learning

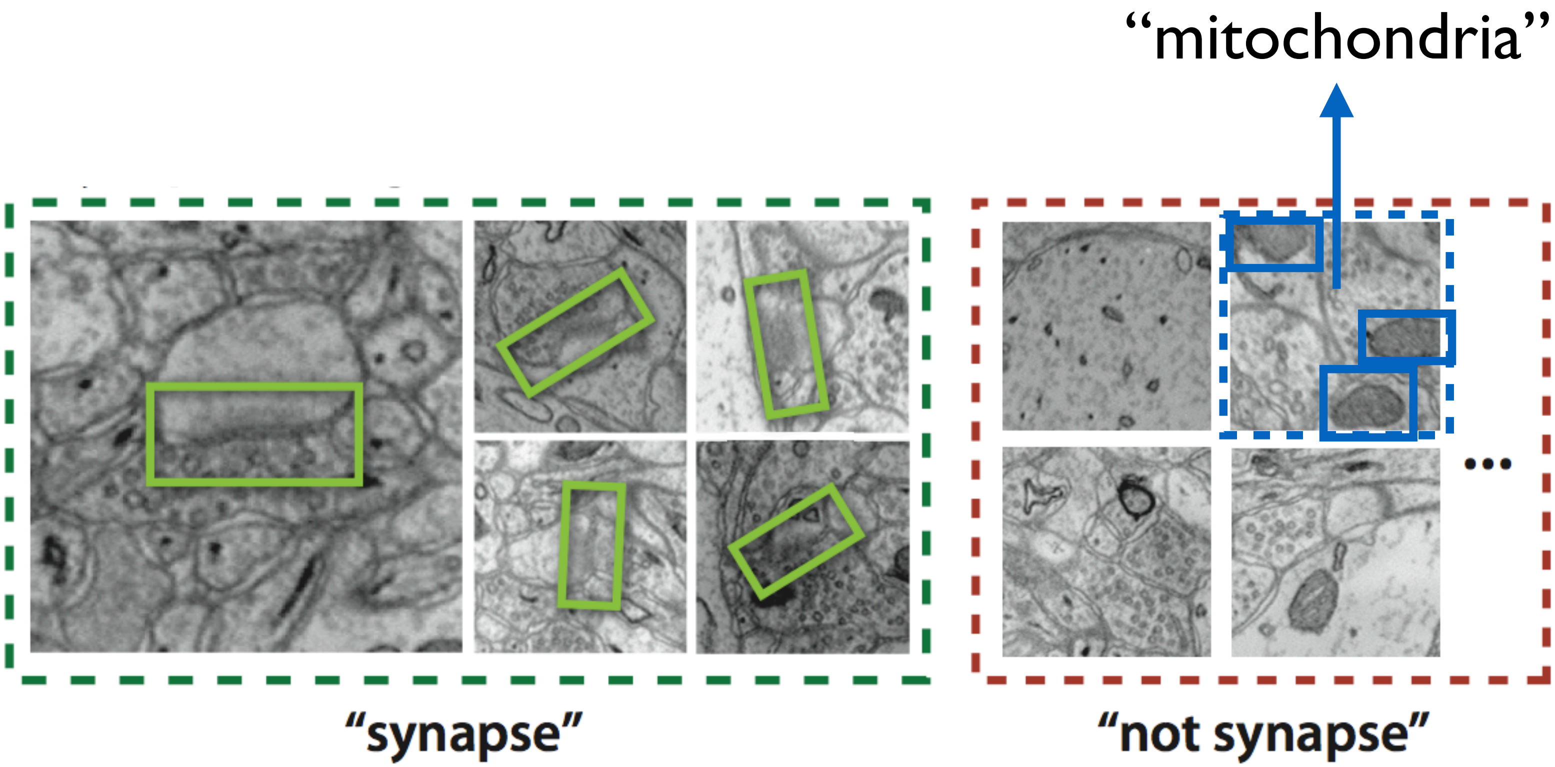

Extreme Value Machine (EVM)

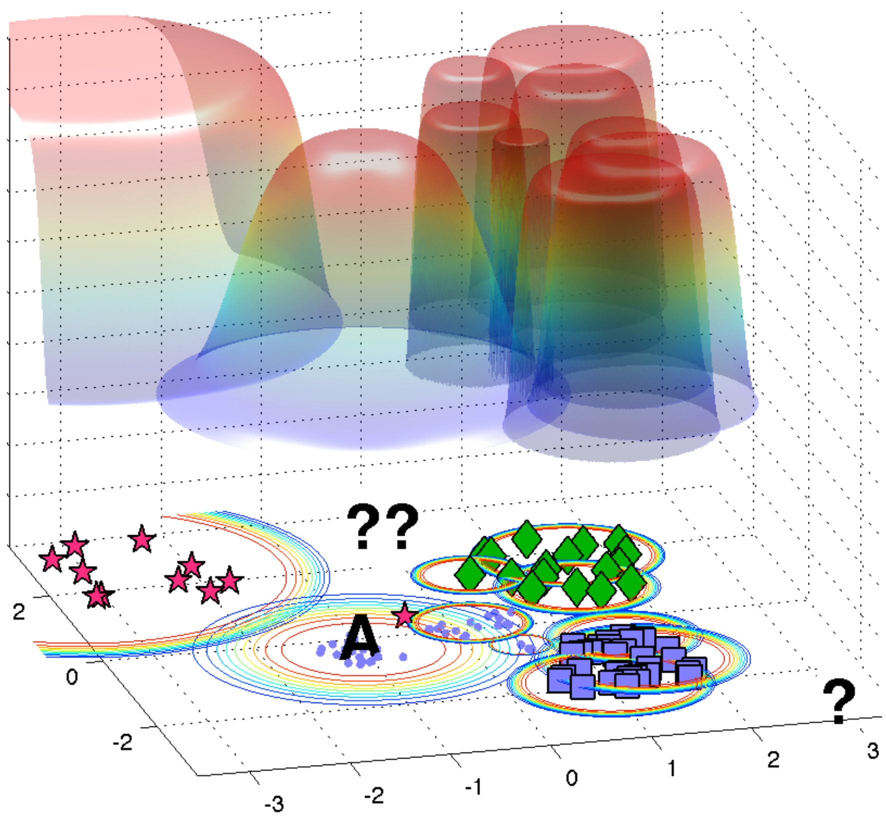

EVM Model for four known classes

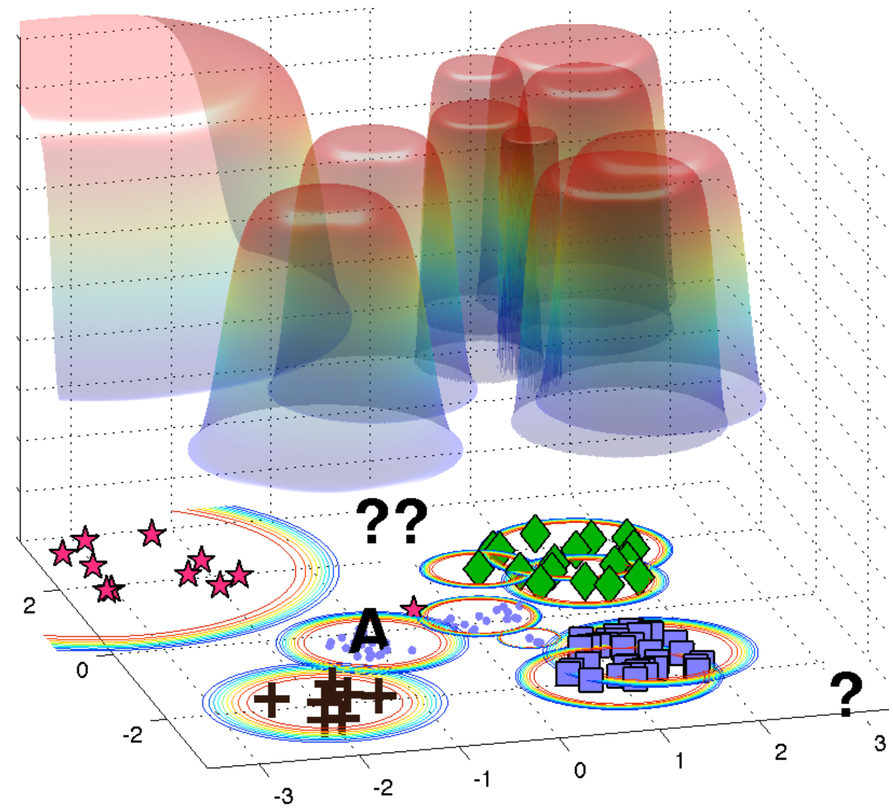

Addition of a new class

Rudd et al. Under Review 2016

EVM Learning Objective:

Number of Extreme Vectors Retained

$$\text{minimize} \sum_{i=1}^{N_l} I(x_i) \text{ subject to}$$

$$N_l = \sum_{i=1}^{N_l} \left[\left(\sum_{j=1}^{N_l} I(x_i) \hat{\Psi}(x_i, x_j, \hat{\kappa}_i, \hat{\lambda}_i) \right) \geq \varsigma \right]$$

Vector Pair and EVT Model for x_i

Probability Threshold for Redundancy

Incremental learning: ImageNet

*Nearest Non-Outlier

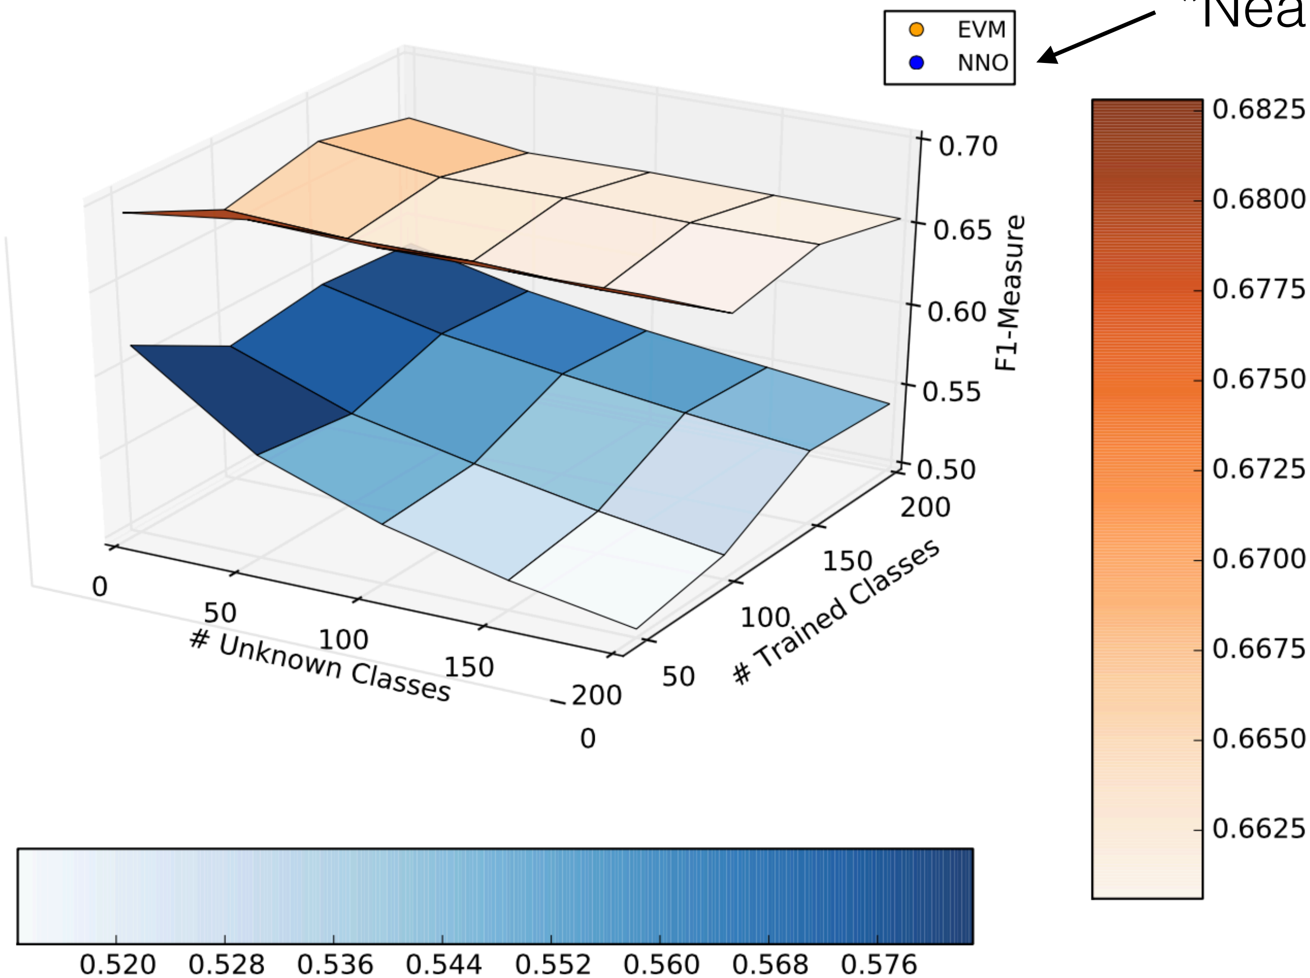

Bendale and Boult CVPR 2015

Basic Machine Learning Benchmark: LETTER

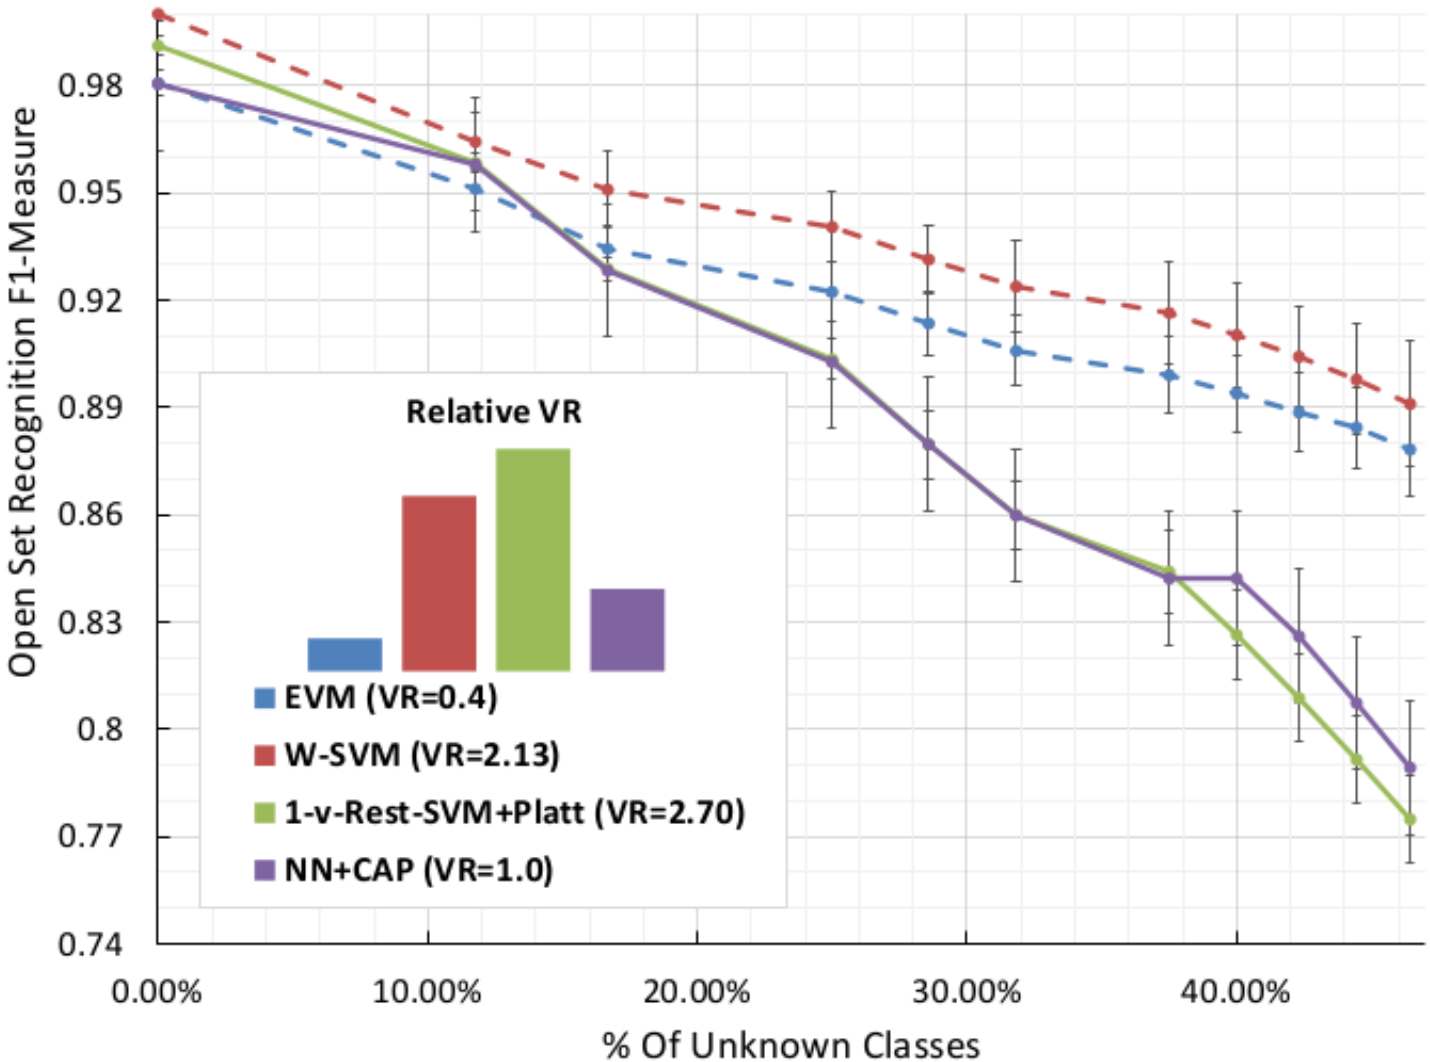

Thank you!

(web) www.wjscheirer.com
(code coming soon) https://github.com/CVRL