Data science in cell imaging
Lecture 5: deep learning in microscopy

"The Great Wave off Kanagawa", by Hokusai, ~1830 (Source: Wikipedia)

All slides are open under the cc-by license

You are free to share and adapt any content from this presentation provided that you attribute the work to its author and respect the rights and licenses associated with its components

PPTX slides available here

Today and next week

- (Discussing potential course projects)
- Enhancing cell image quality with deep learning
- Generative models for cell structure with deep learning
- Classifying cell state with deep learning

Machine learning 101

Recommended deck for non-computational students

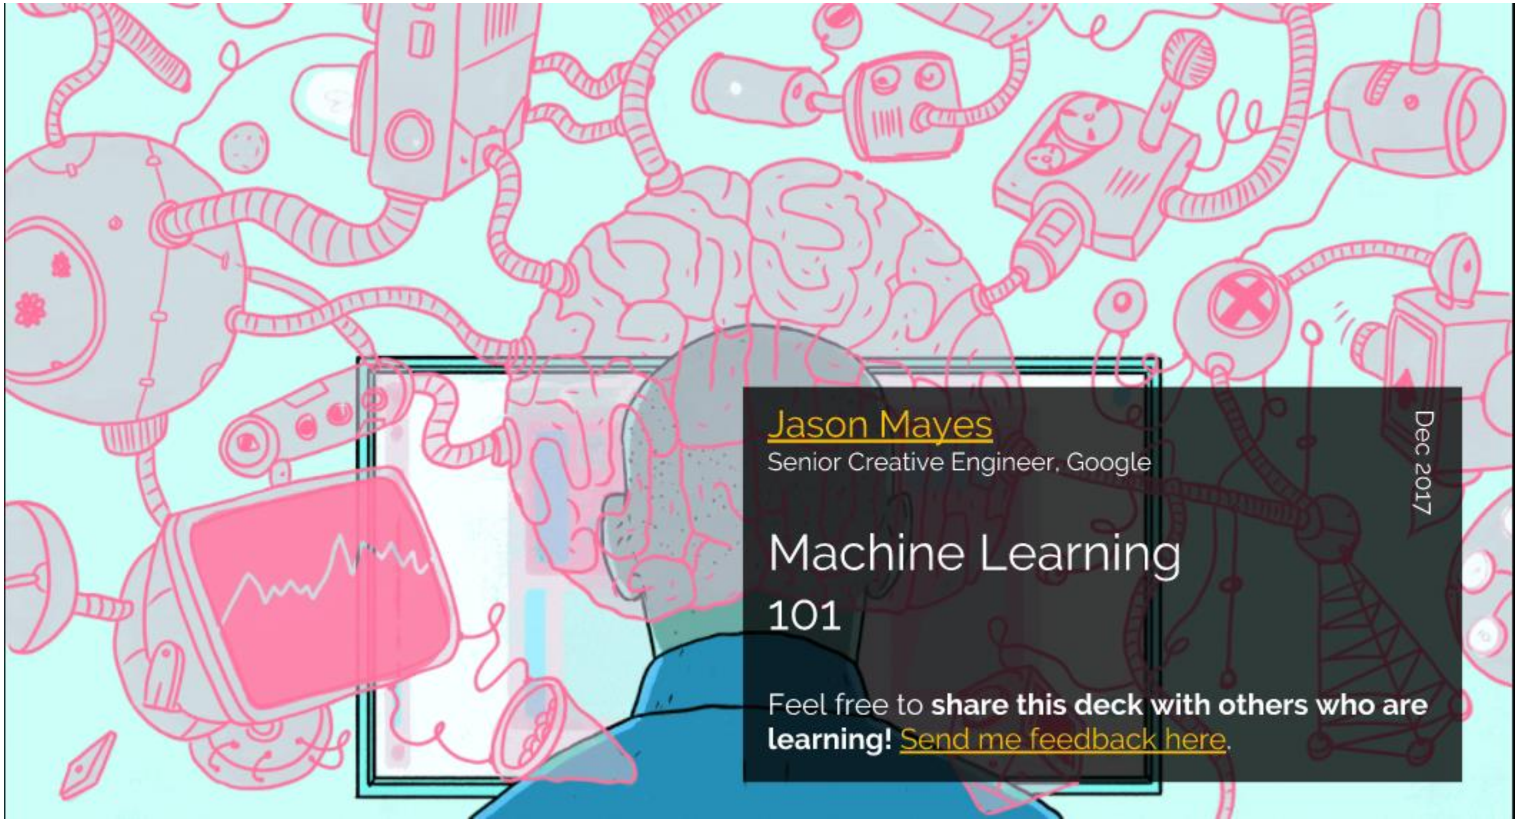

Supervised machine learning

Data → feature extraction → training → discriminative model

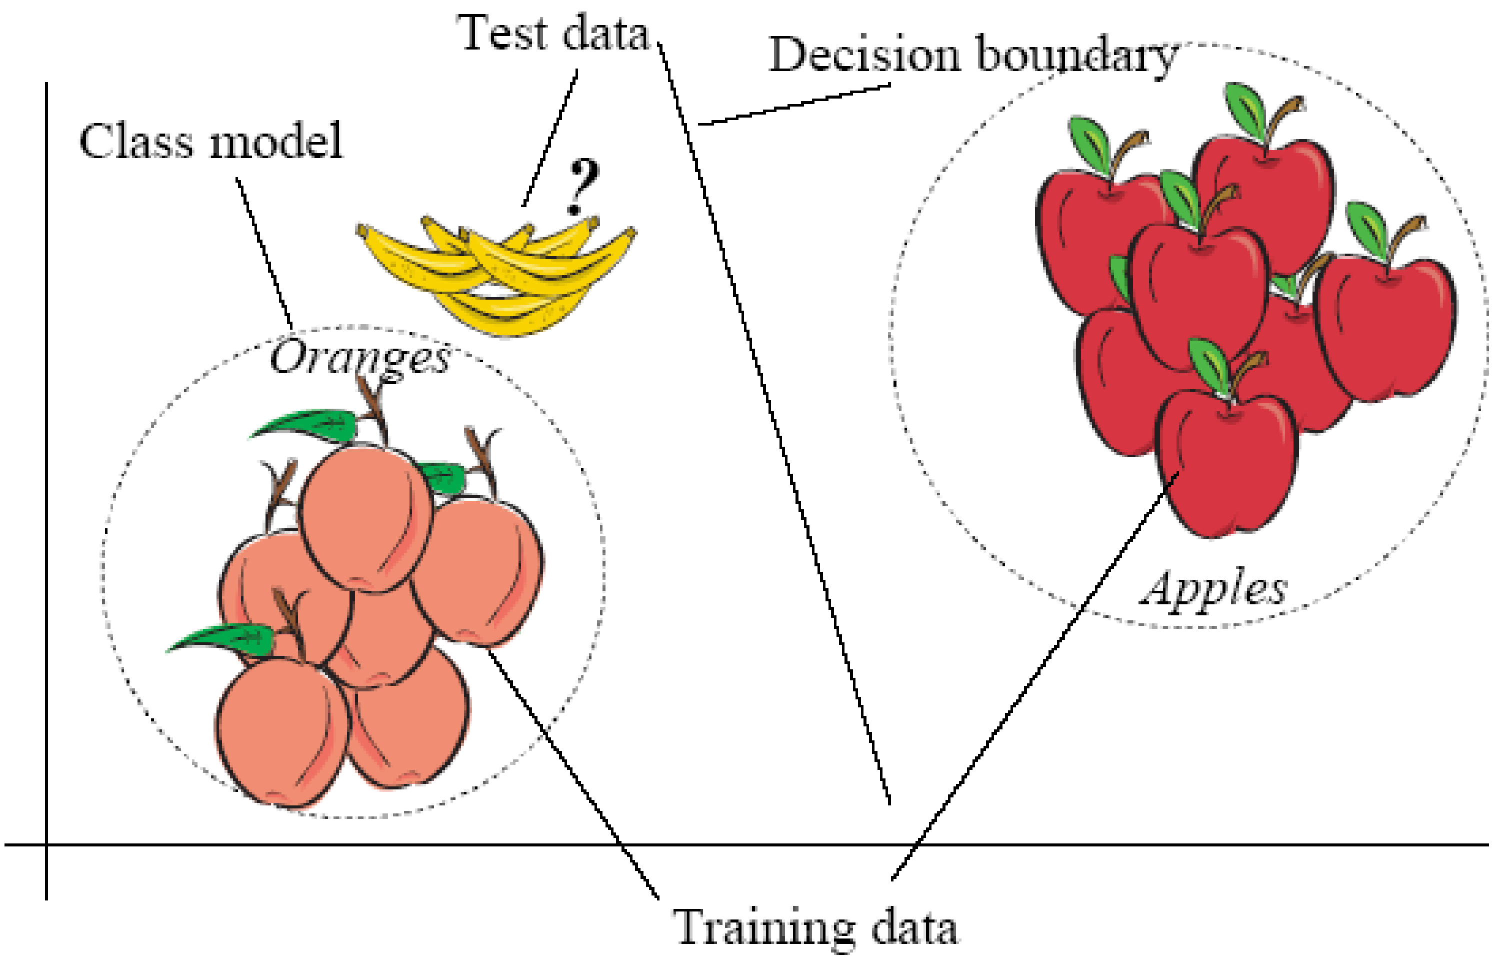

Images by MIT OCW.

Credit: B. Heisele, Y. Ivanov, T. Poggio

Unsupervised machine learning

Data → feature extraction → clustering

Source: https://bit.ly/2IJN0NT

Deep learning

A powerful integration of automated feature extraction and model training

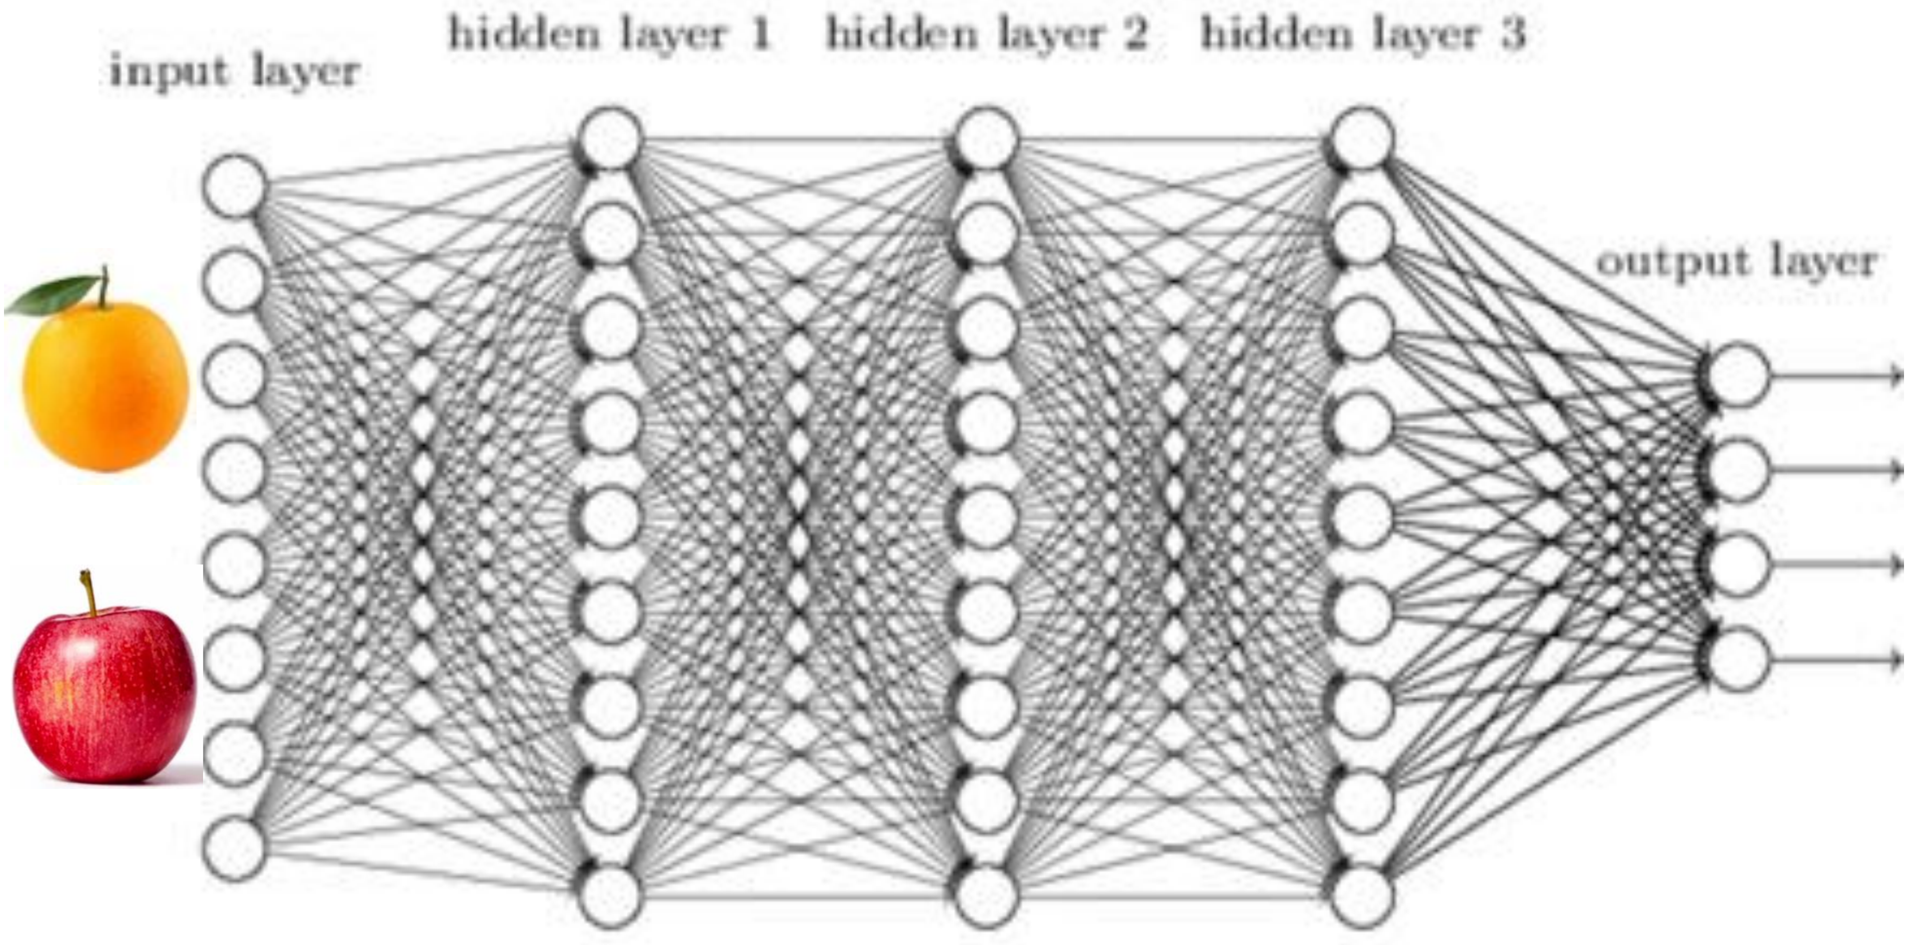

Adapted from https://bit.ly/2GMZou6

How does it works?

Artificial neuron (perceptron)

A neuron simply has **a bunch of weighted inputs** (take the input number and multiply by the weight) **that are summed together.** A bias is then added to this total. The weights and bias are determined when we train the system. If final result is greater than a threshold, it activates, providing an output. Strength of output depends on the activation function chosen. The output is then fed into other neurons and the process repeats.

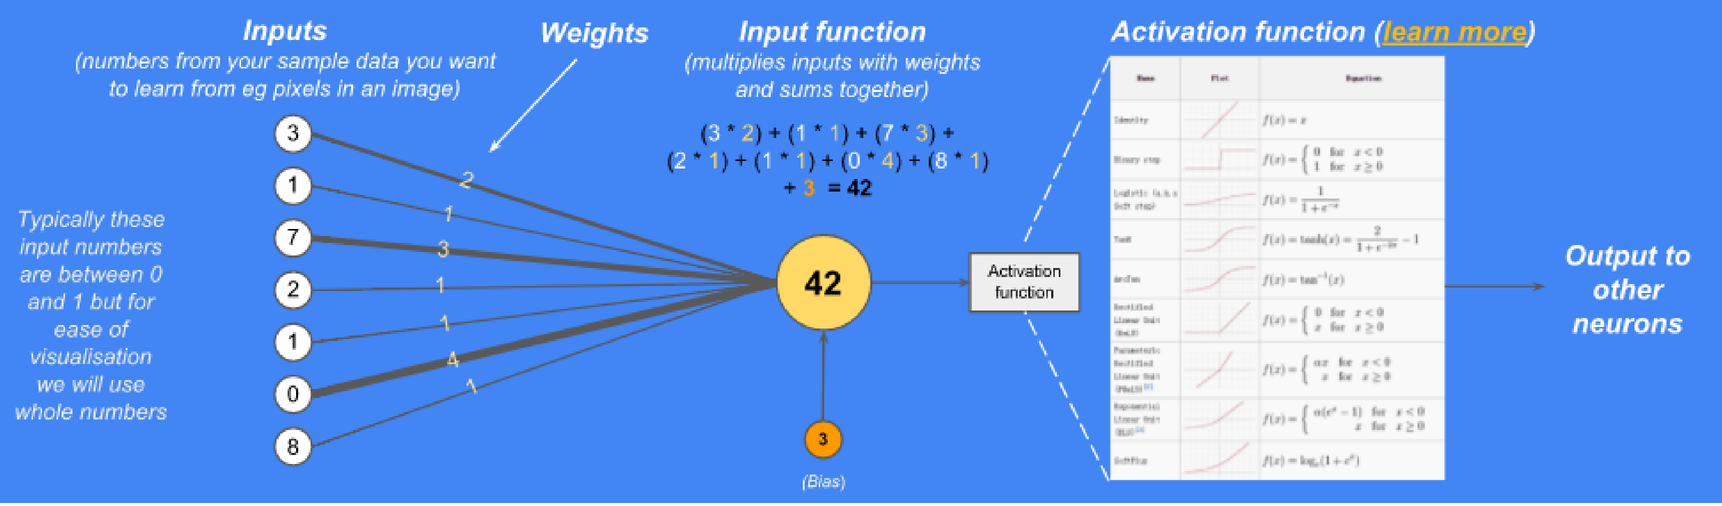

Machine Learning 101 by Jsaon Mayes, http://tiny.cc/ly9cnz

Deep neural network

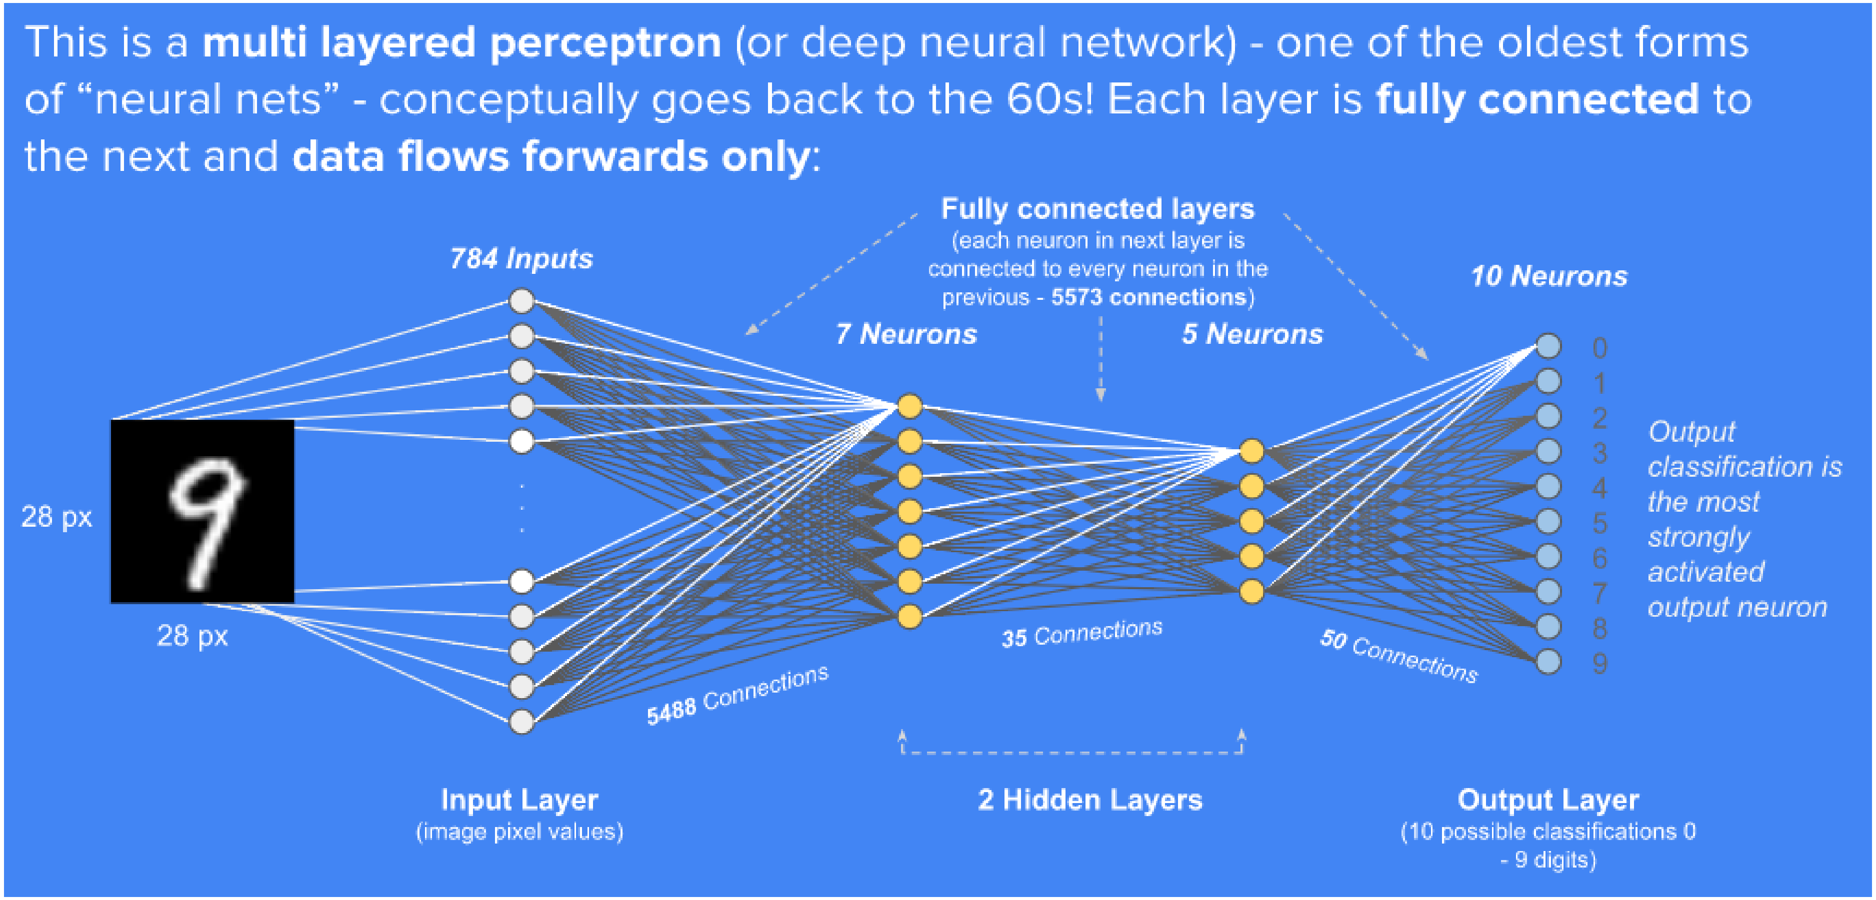

Machine Learning 101 by Jsaon Mayes, http://tiny.cc/ly9cnz

Deep neural network

A Deep Neural Network (DNN) simply consists of many "hidden layers" between the input and the output. Each layer can learn from the one before it from which higher level learning can take place. These hidden layers typically are of lower dimensionality so they can generalise better and not overfit to the input data. These middle layers in the system can learn features of features. For example bunches of "edges" can lead to "face parts" which lead to "faces" that the system can then recognize.

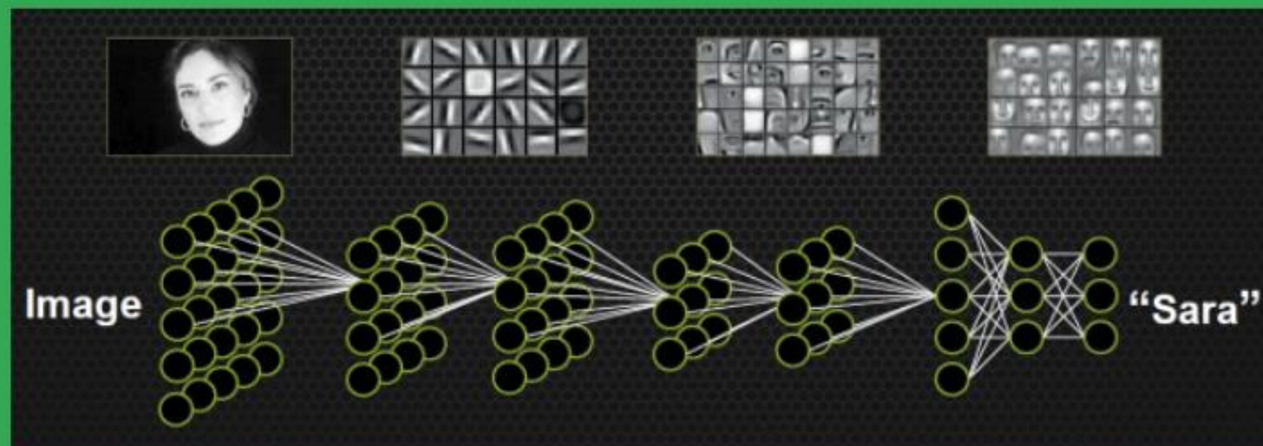

Machine Learning 101 by Jsaon Mayes, http://tiny.cc/ly9cnz

Autoencoder: unsupervised data encoding

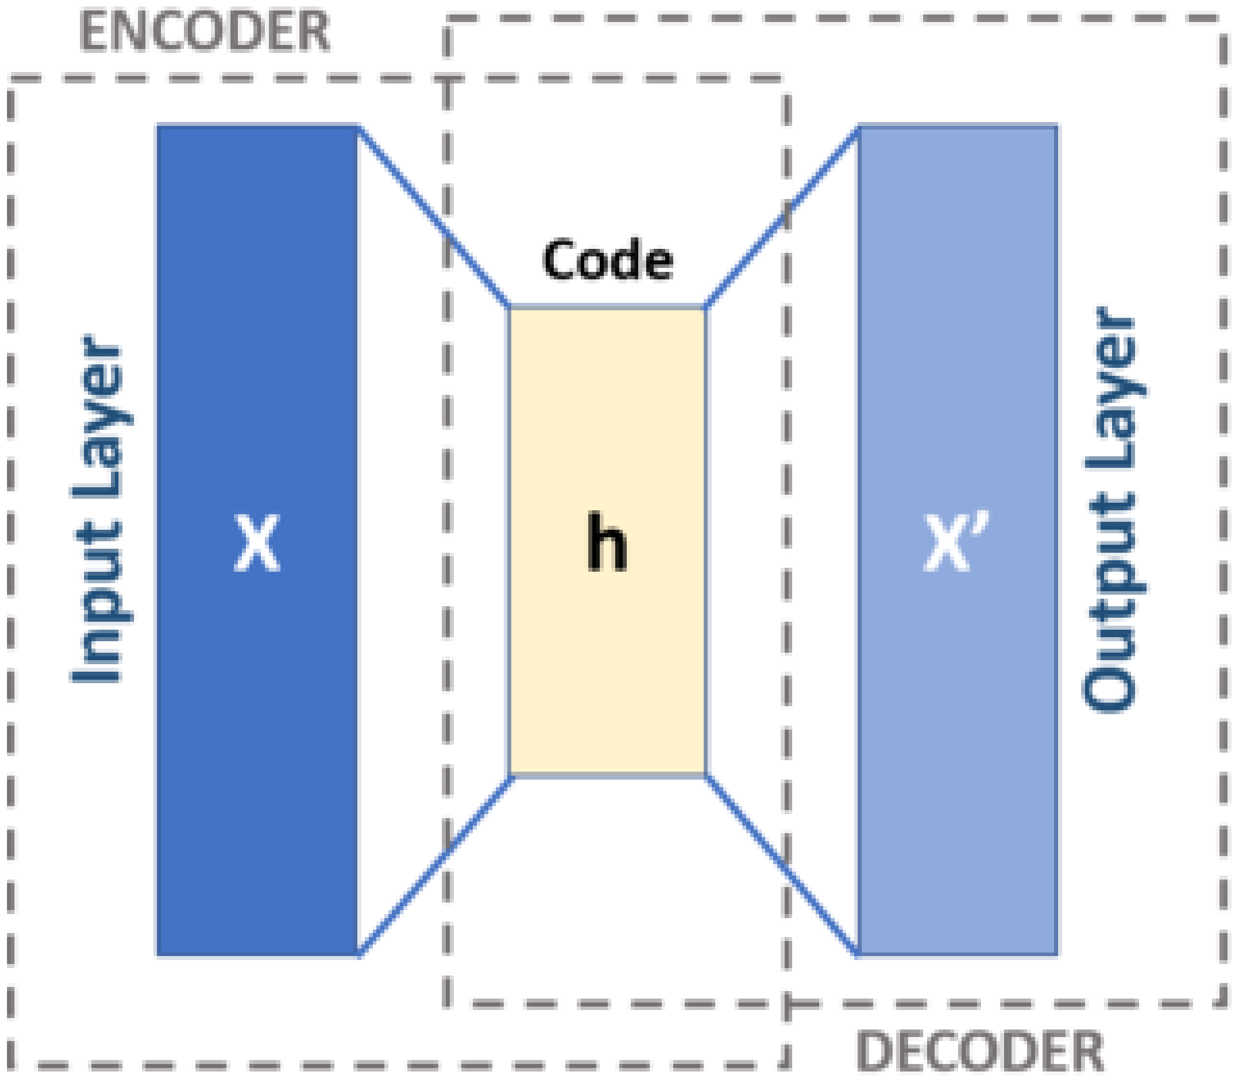

Wikipedia

Many applications for ML in cell biology!

Today I'll review a few specific examples on deep learning for (1) image restoration; (2) prediction of (un-labeled) protein localization patterns

Imaging tradeoffs

Spatial/temporal resolution, experiment time, imaging depth, fluorophore density, bleaching, and photo-toxicity

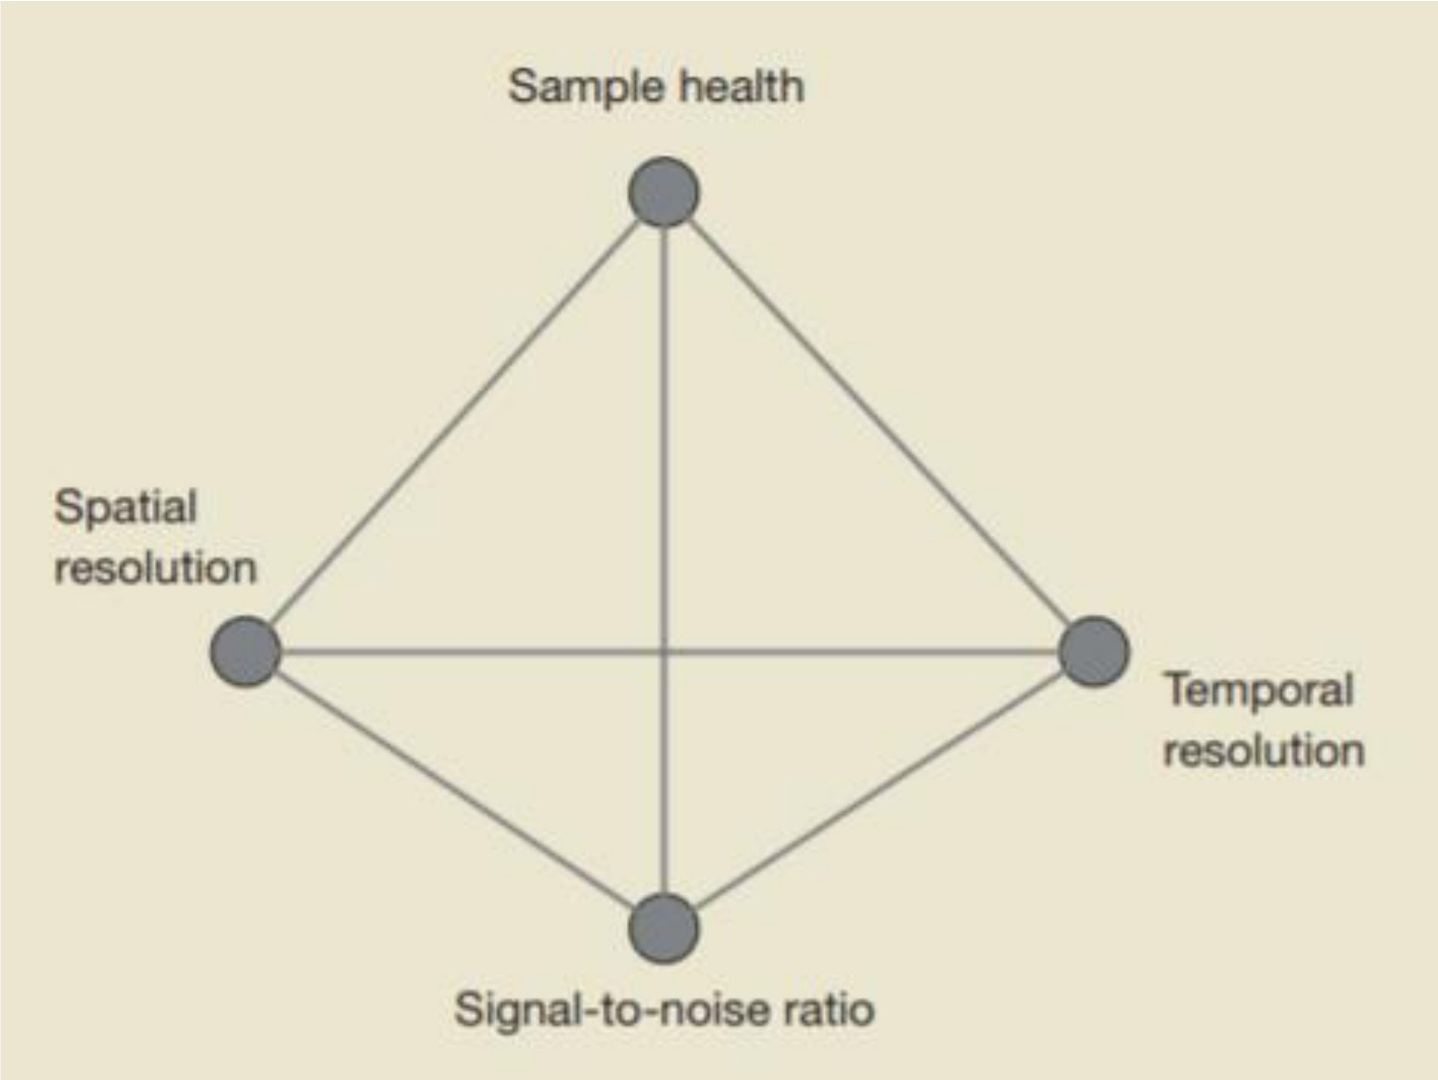

Hadar Aharoni (IDF) from Scherf et al. (2015)

Content aware image restoration (CARE)

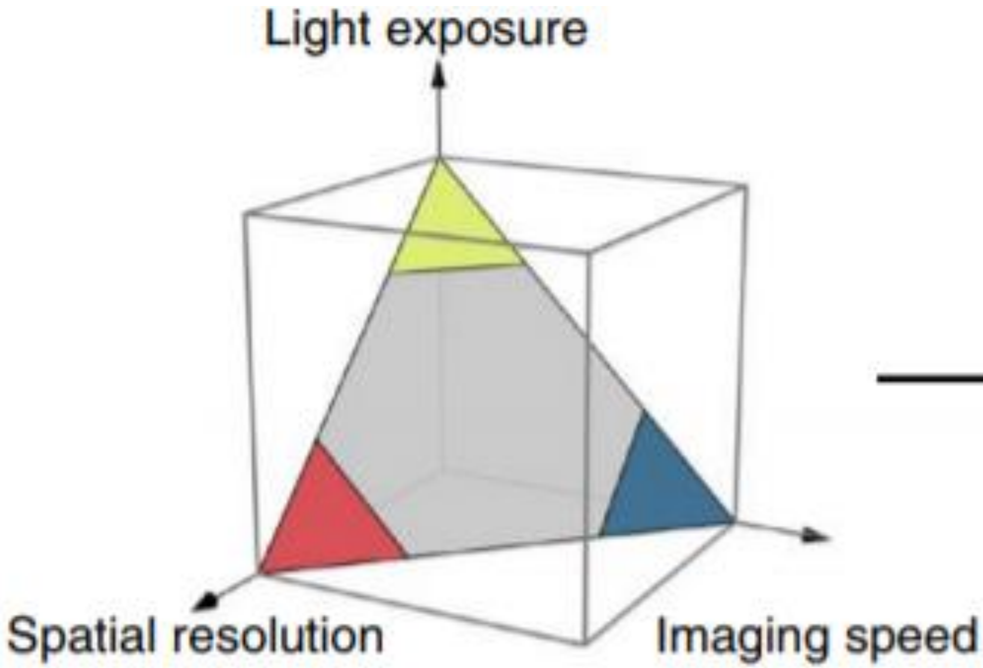

- Traditional methods use general assumptions to perform the restoration
- CARE leverages the available knowledge about the specific experimental task / setup

Weigert et al. (2018)

Content aware image restoration (CARE)

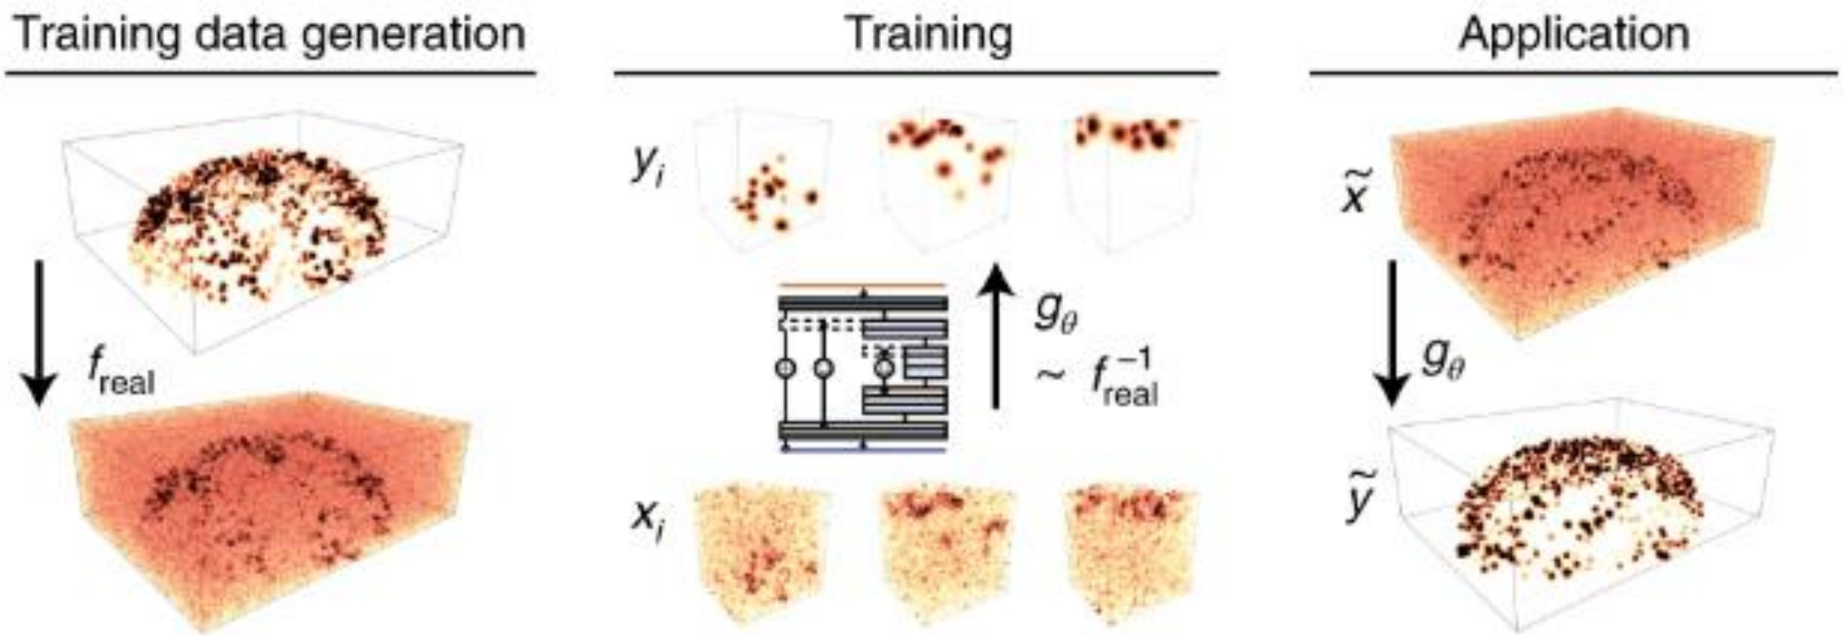

- Traditional methods use general assumptions to perform the restoration
- CARE leverages the available knowledge about the specific experimental task / setup

Weigert et al. (2018)

The machinery: U-Net

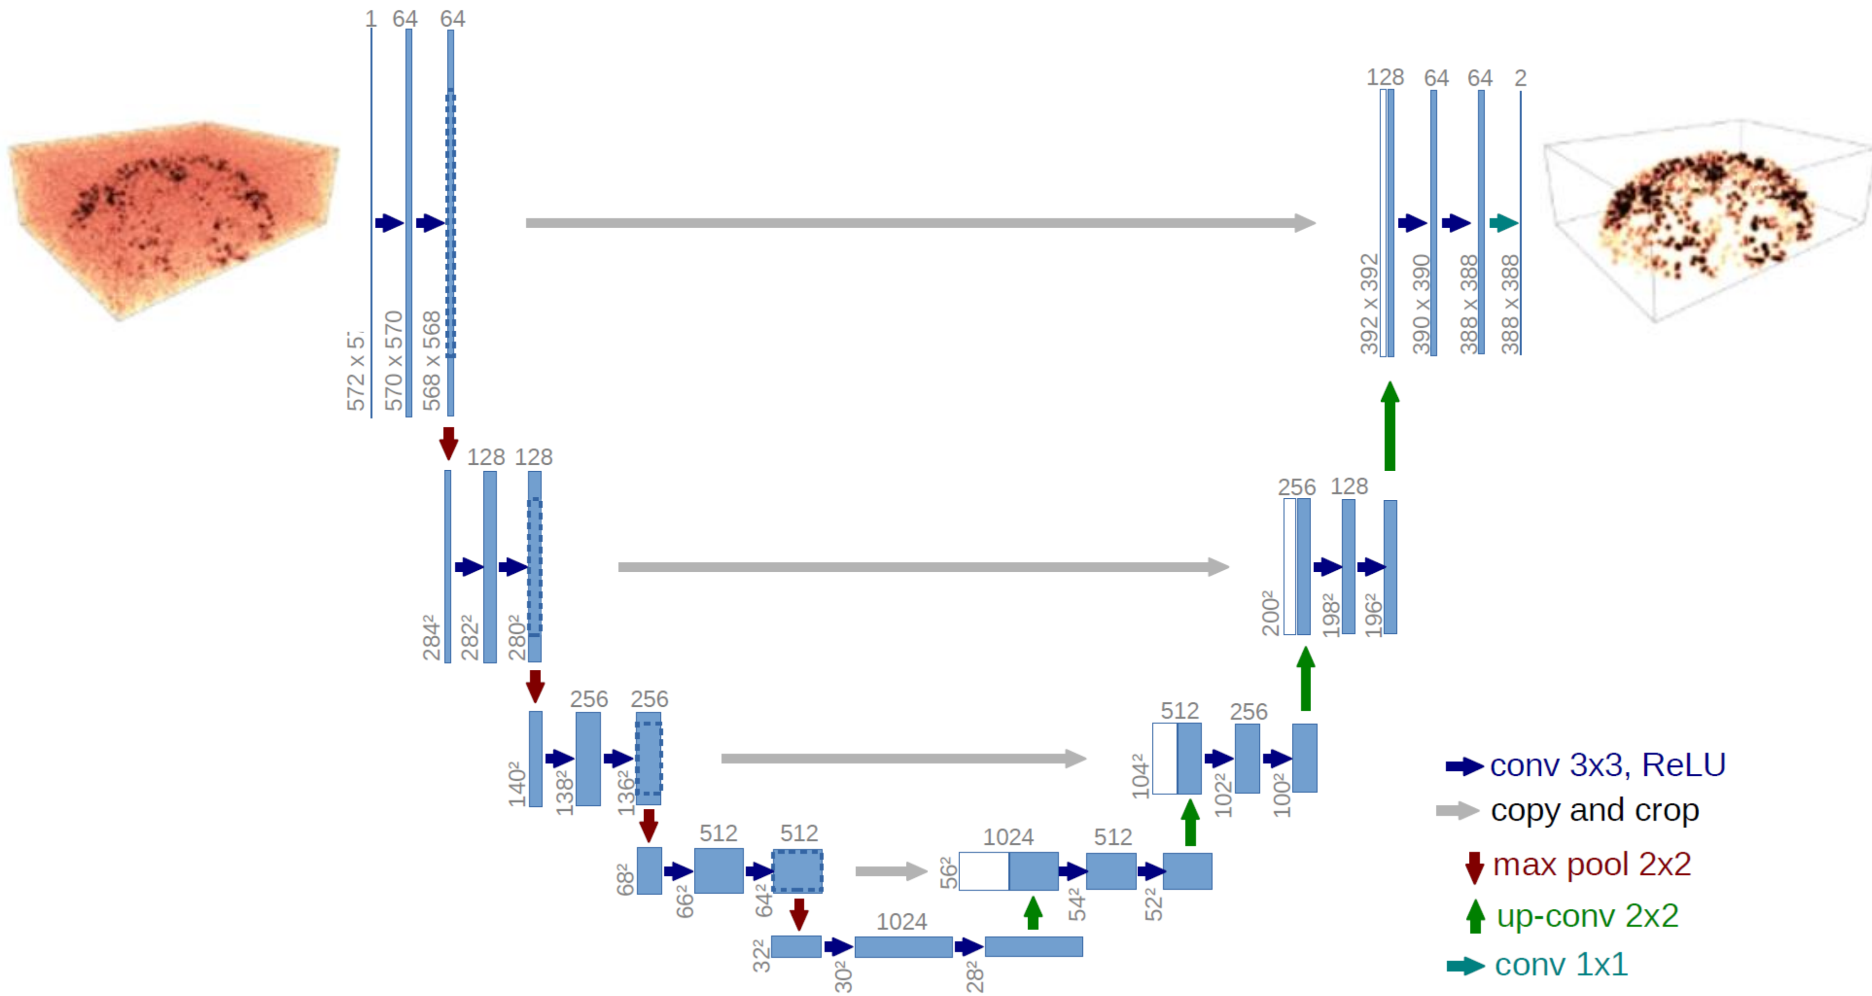

Ronneberger et al. (2015)

Results

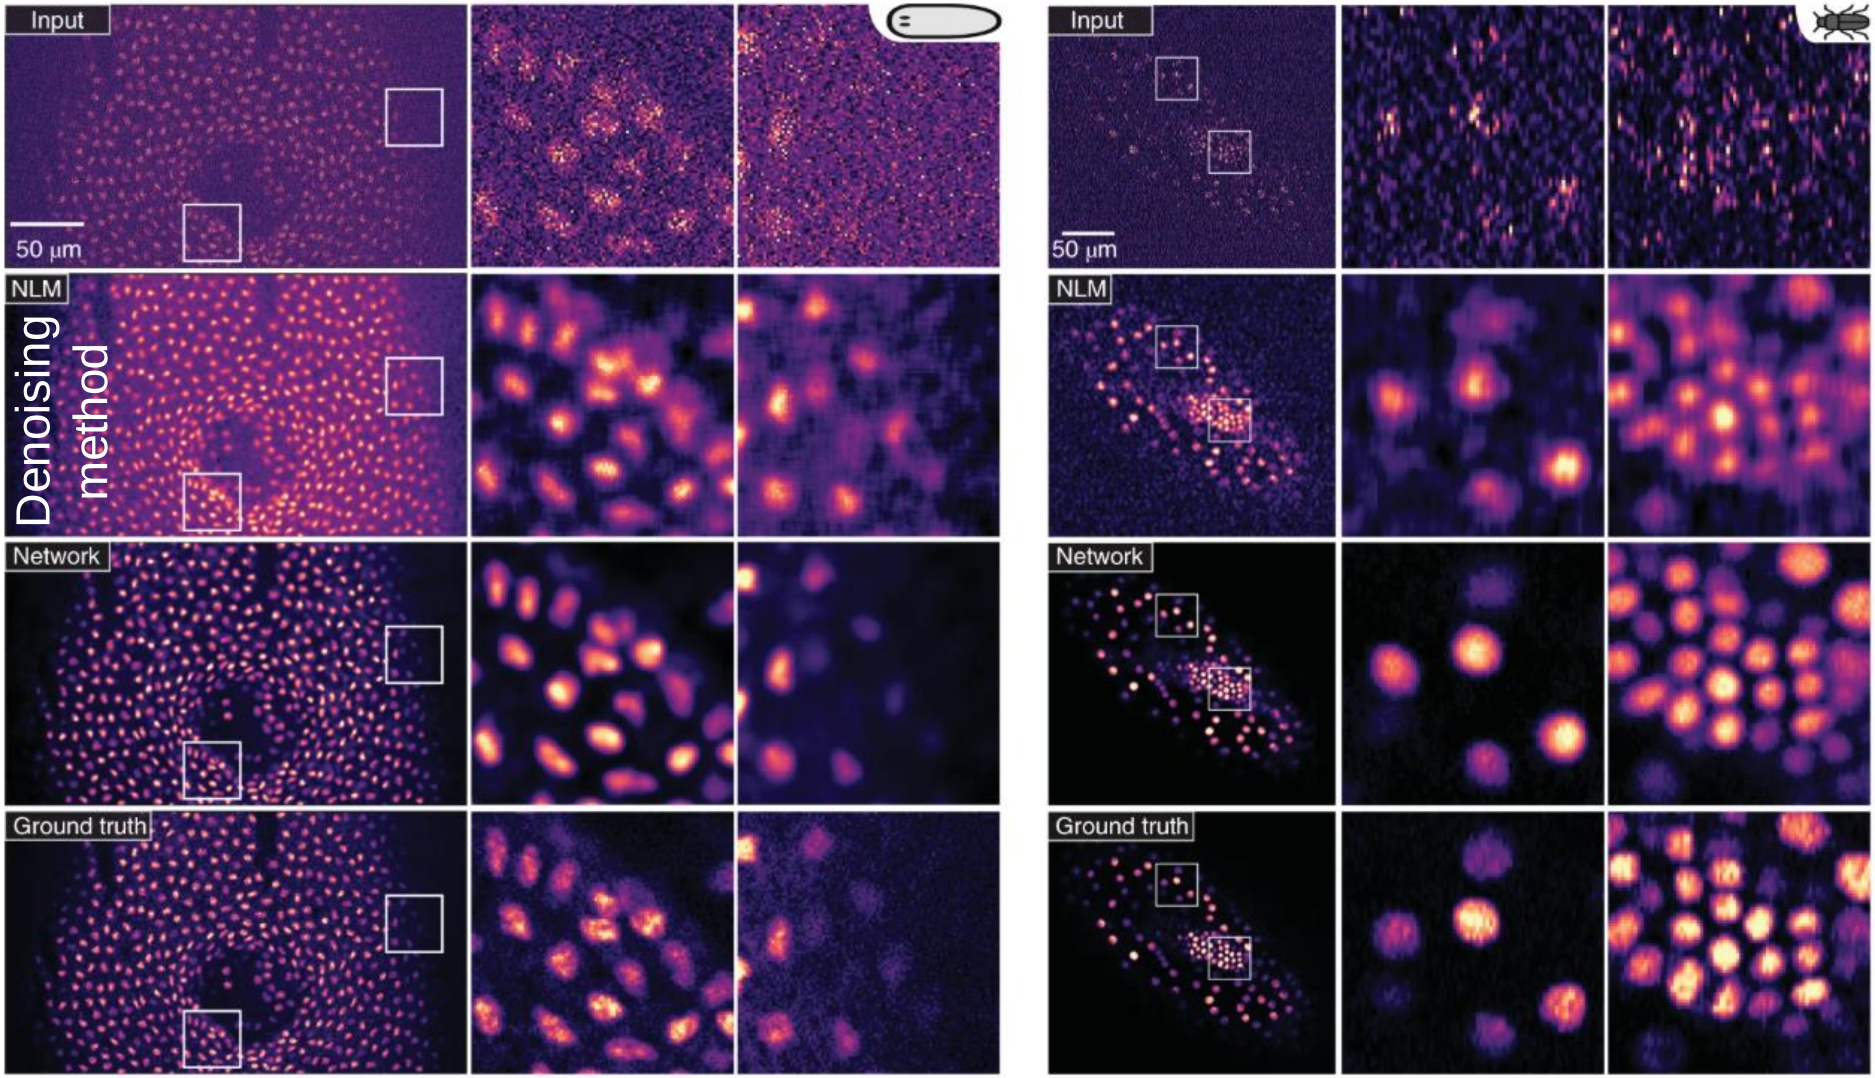

Weigert et al. (2018)

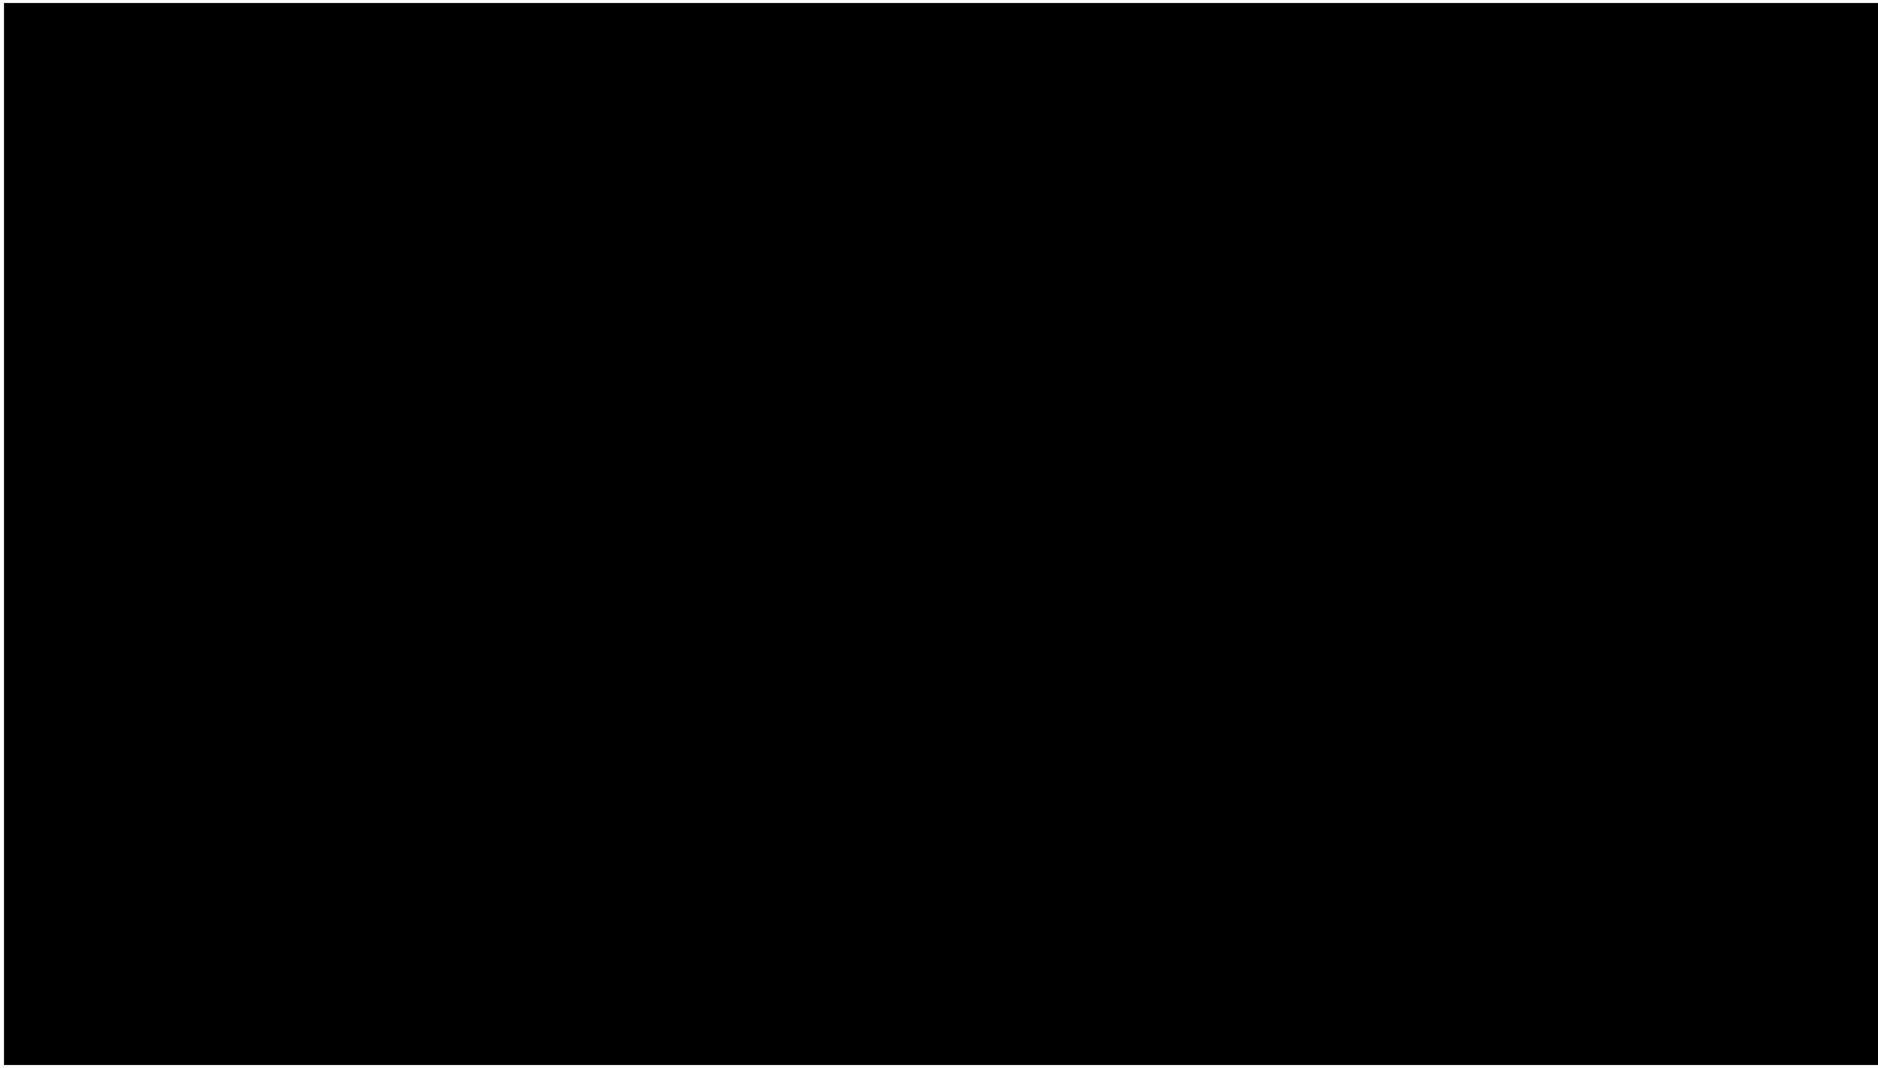

Weigert et al. (2018)

Restoration of low-SNR images

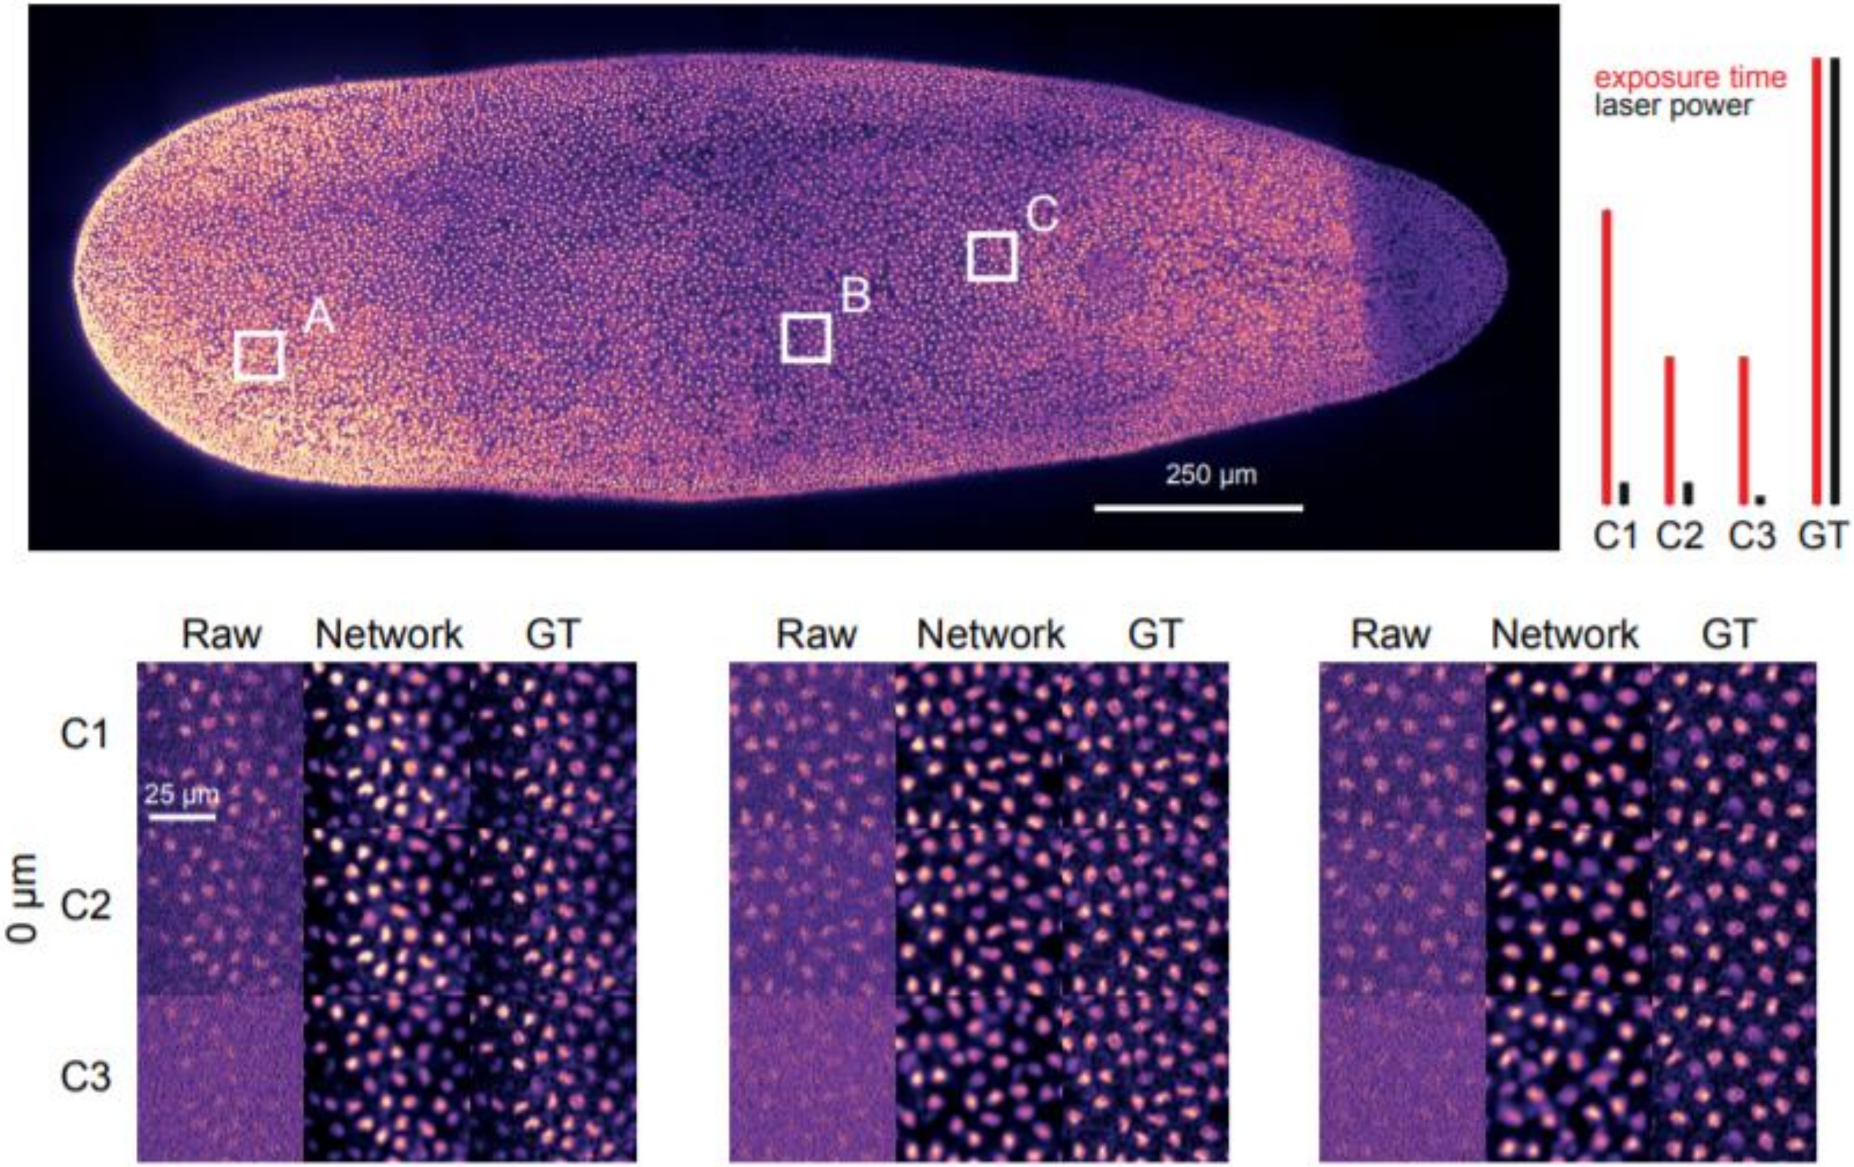

Weigert et al. (2018)

Restoration of low-SNR images

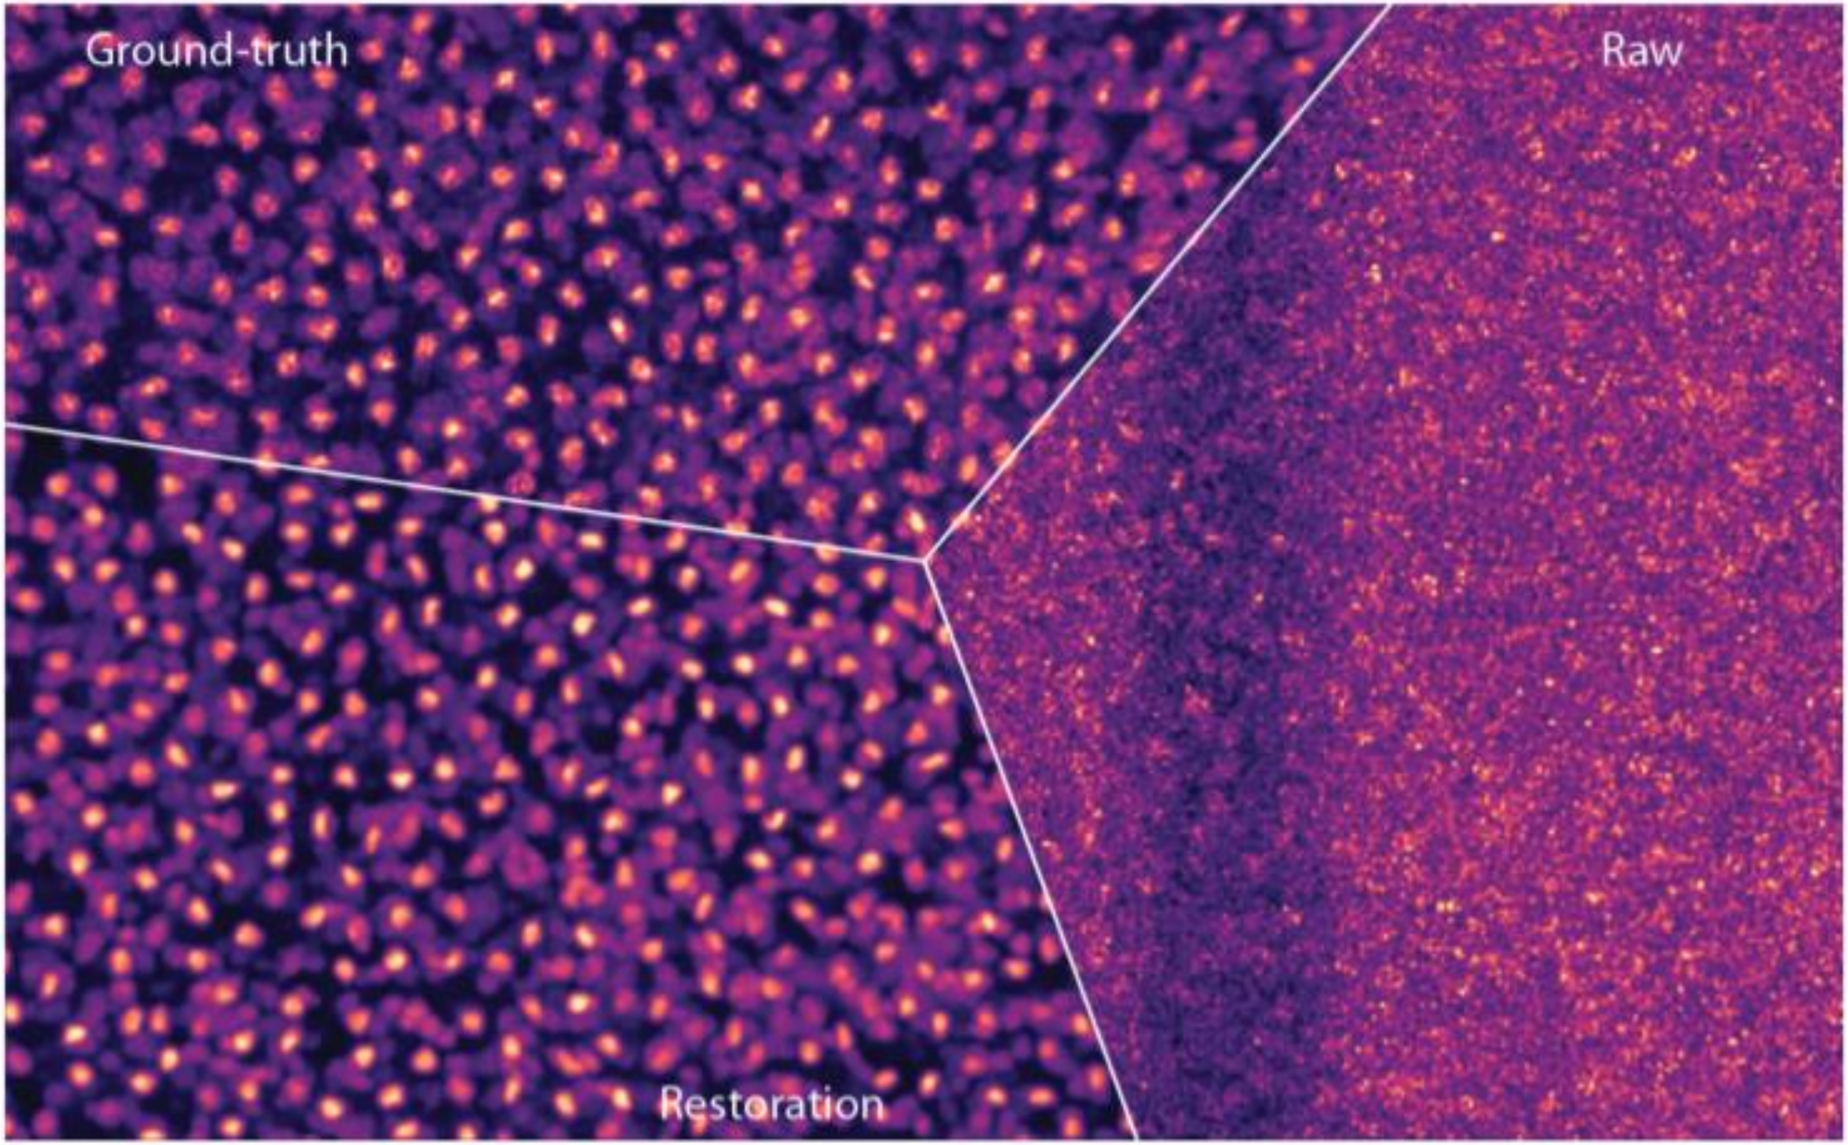

Weigert et al. (2018)

Comparison: different denoising methods

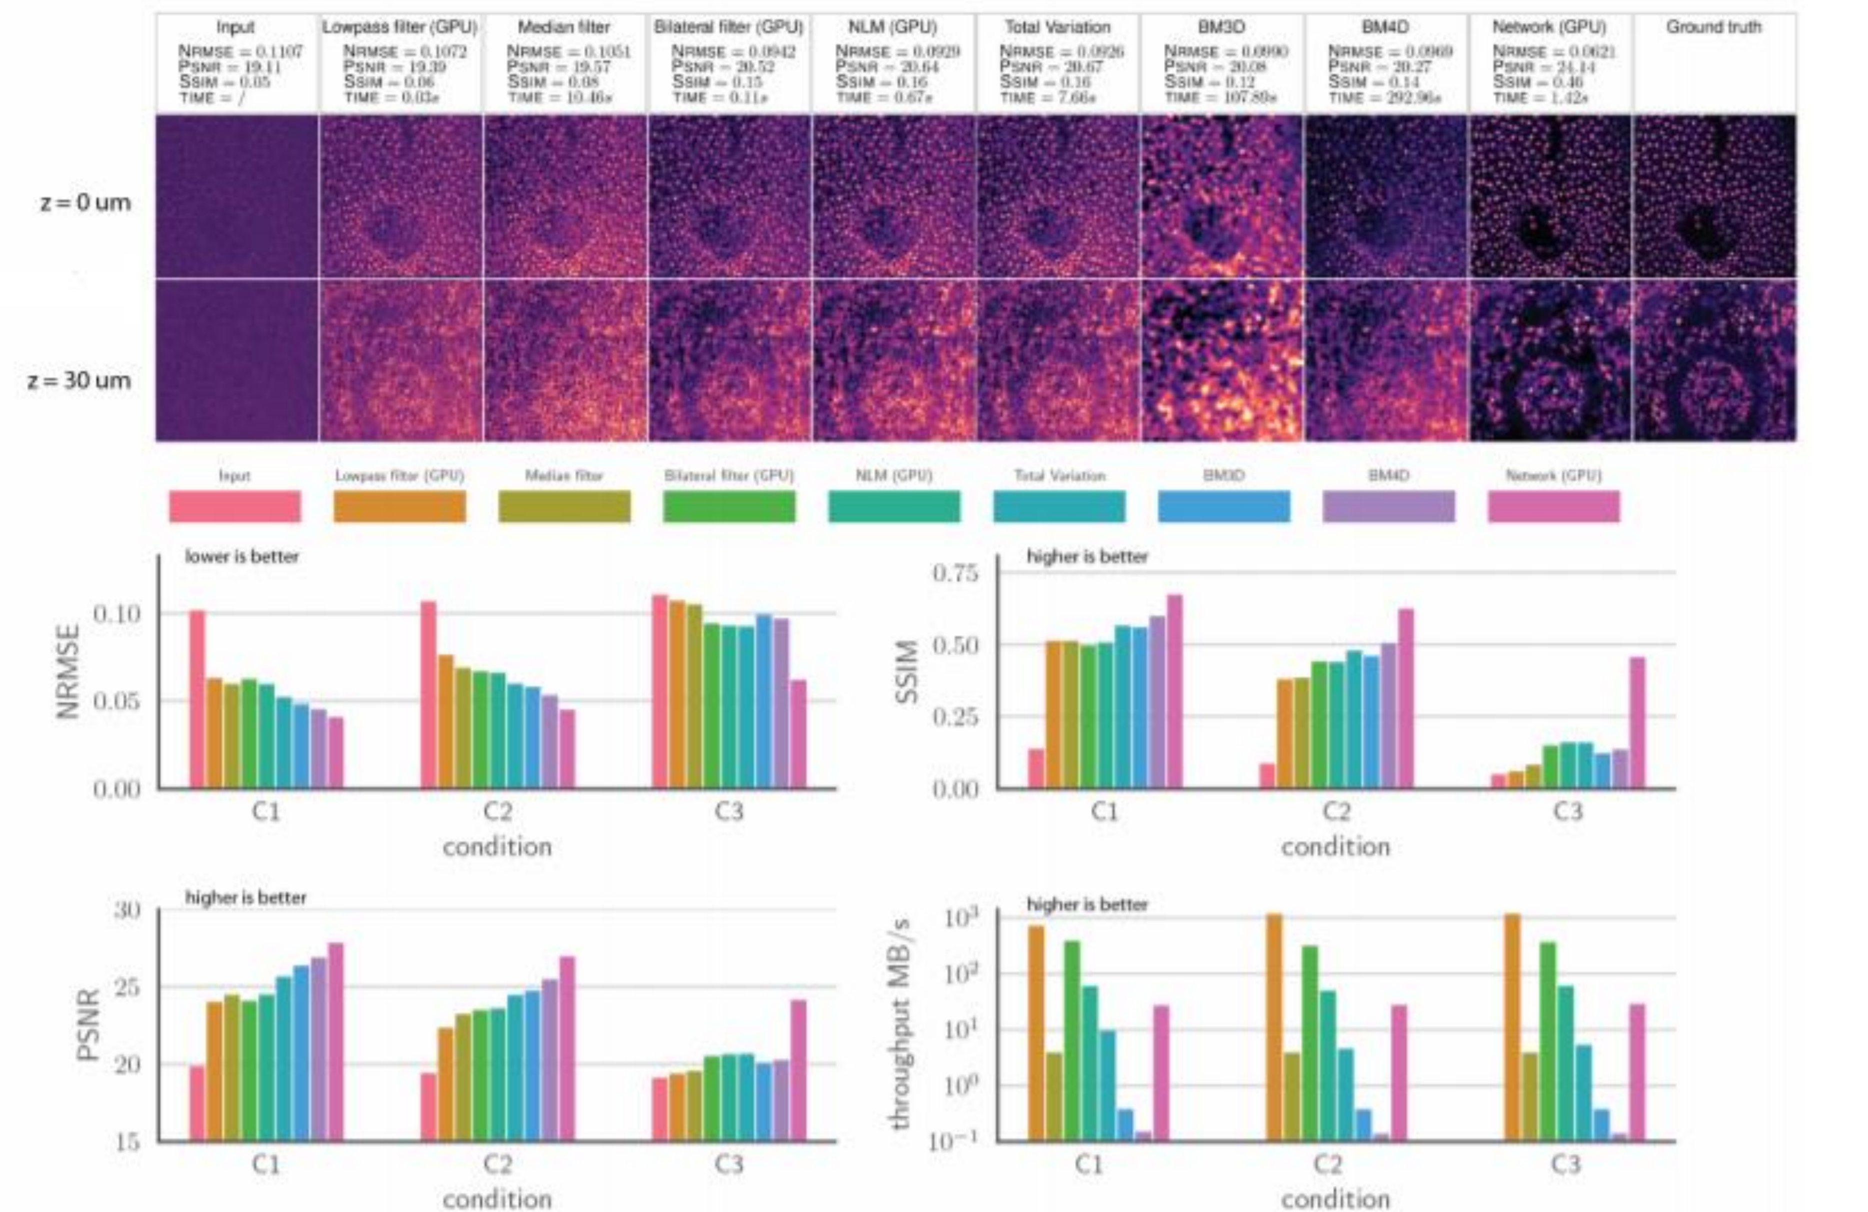

Weigert et al. (2018)

CARE improves downstream analysis

Cell segmentation and tracking of a projected time-lapse

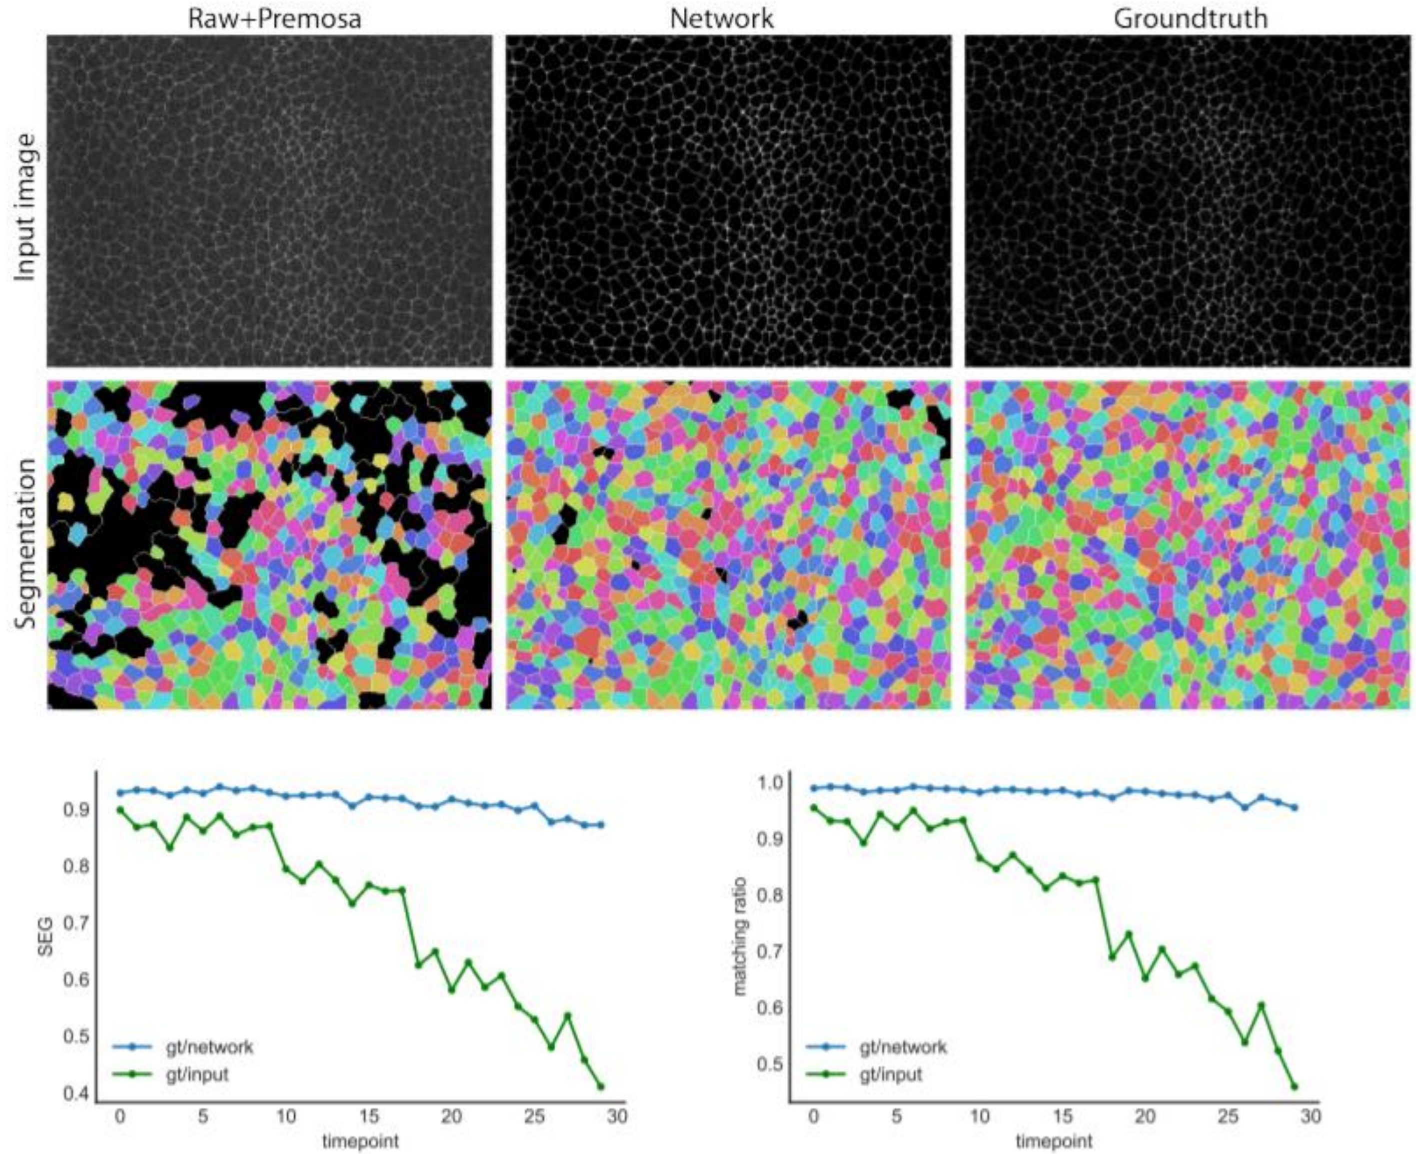

Weigert et al. (2018)

Image restoration of unseen axial slices with semi-synthetic training data

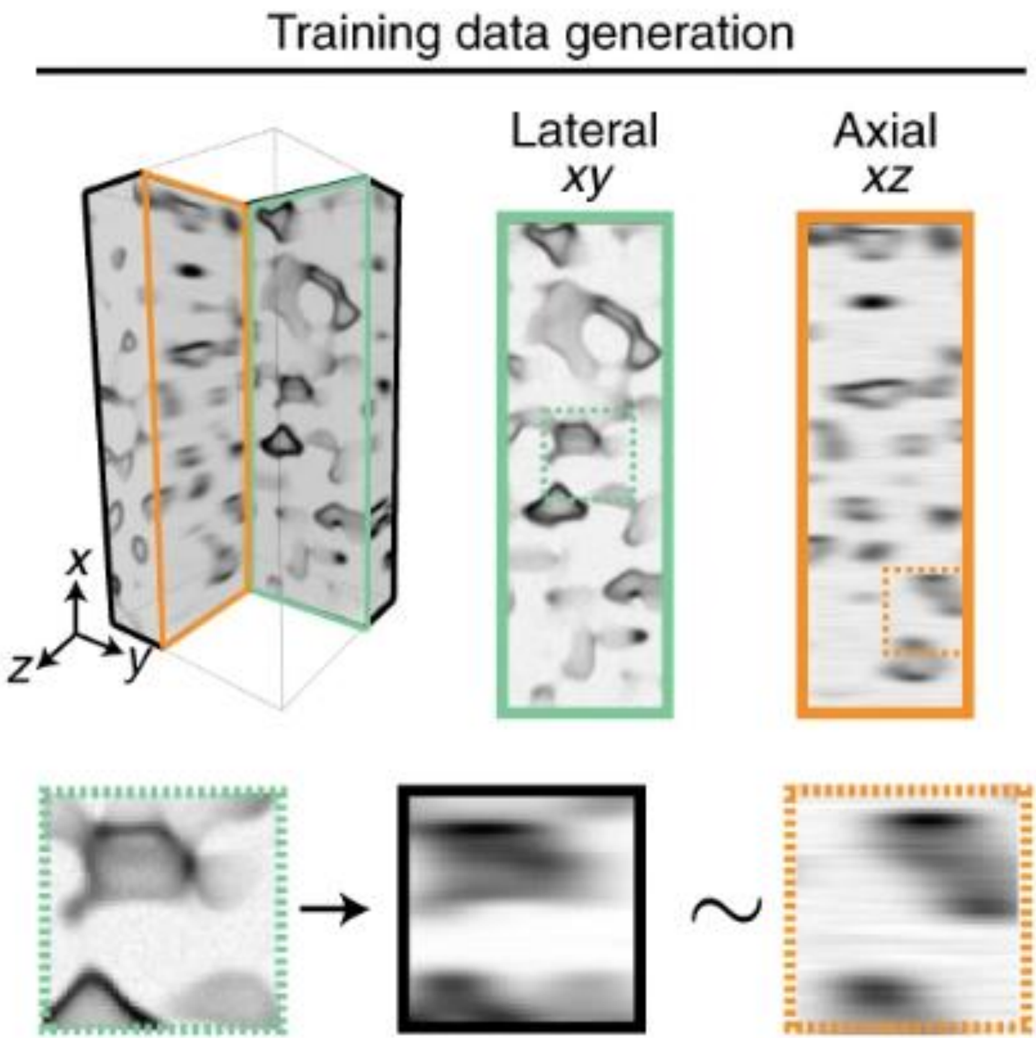

Weigert et al. (2018)

CARE increase axial resolution of acquired volumes

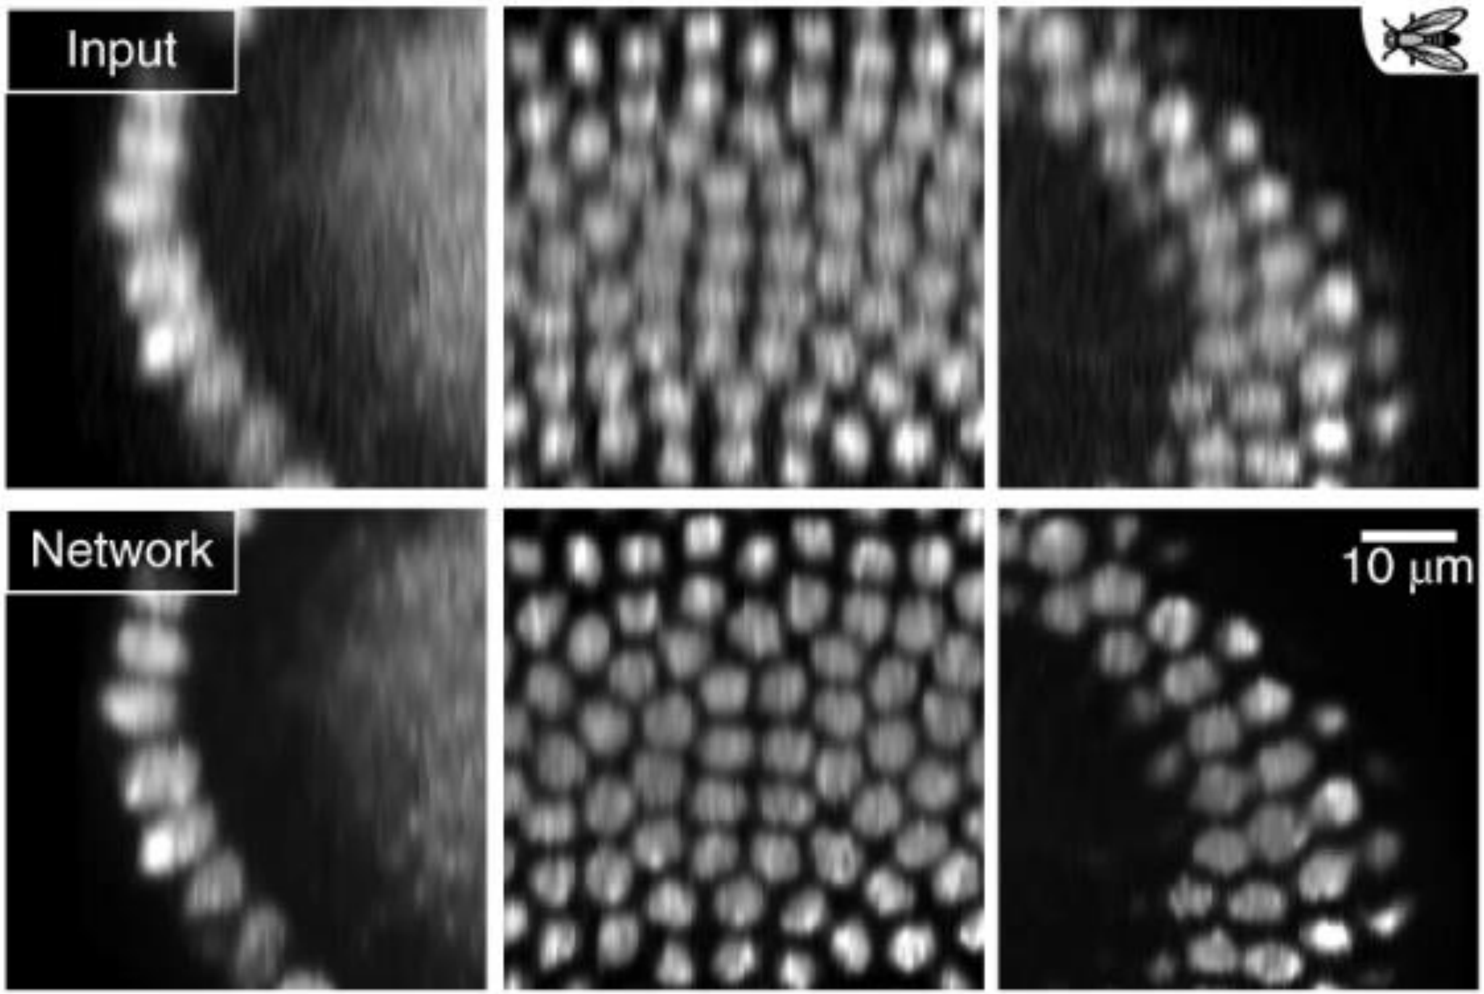

Weigert et al. (2018)

CARE increase axial resolution of acquired volumes

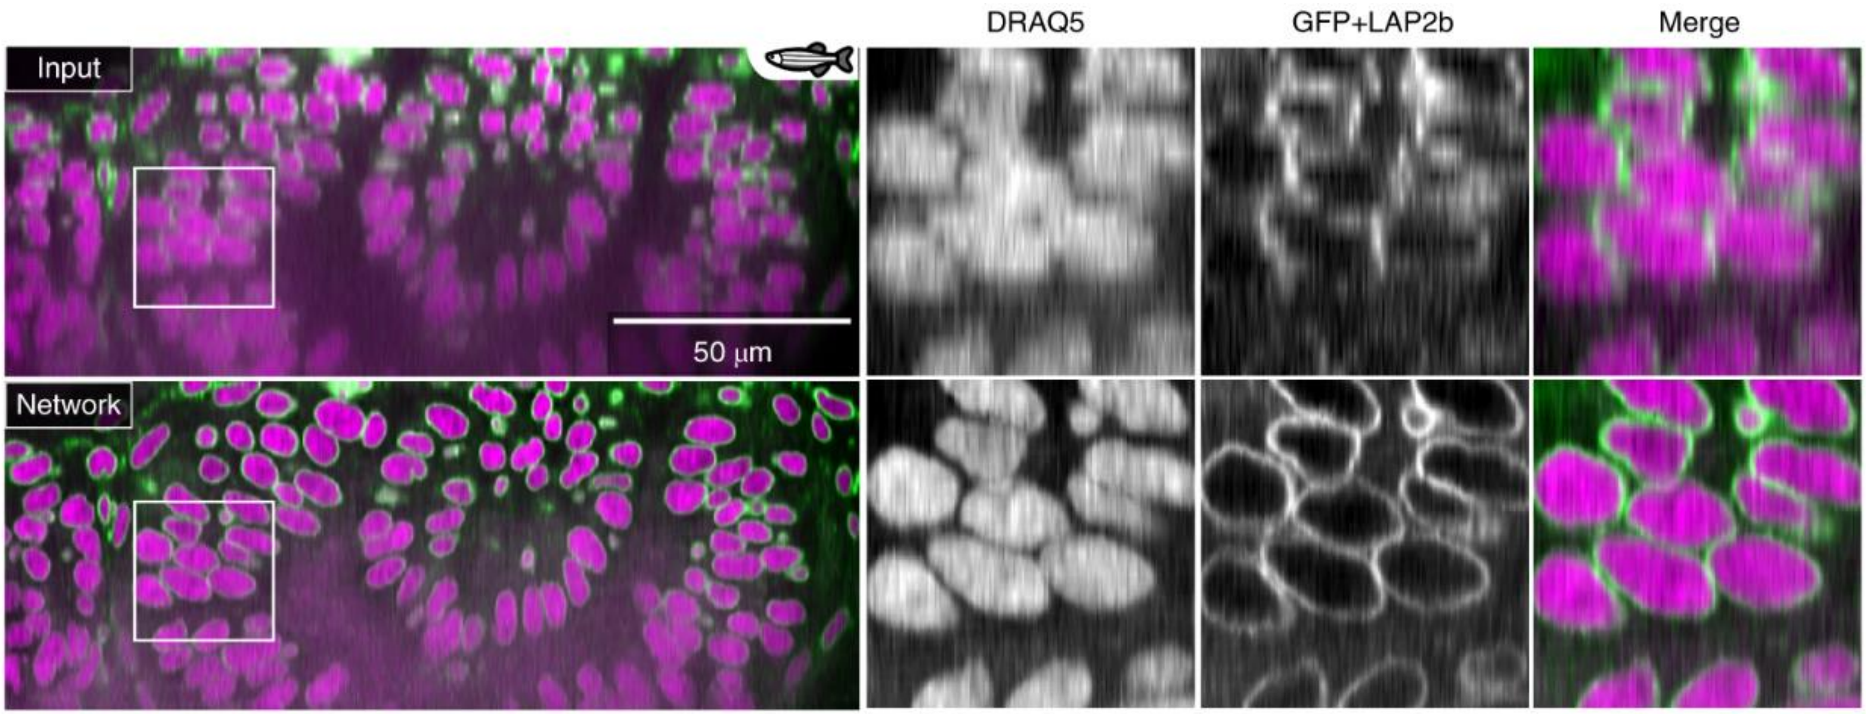

Weigert et al. (2018)

CARE increase axial resolution of acquired volumes

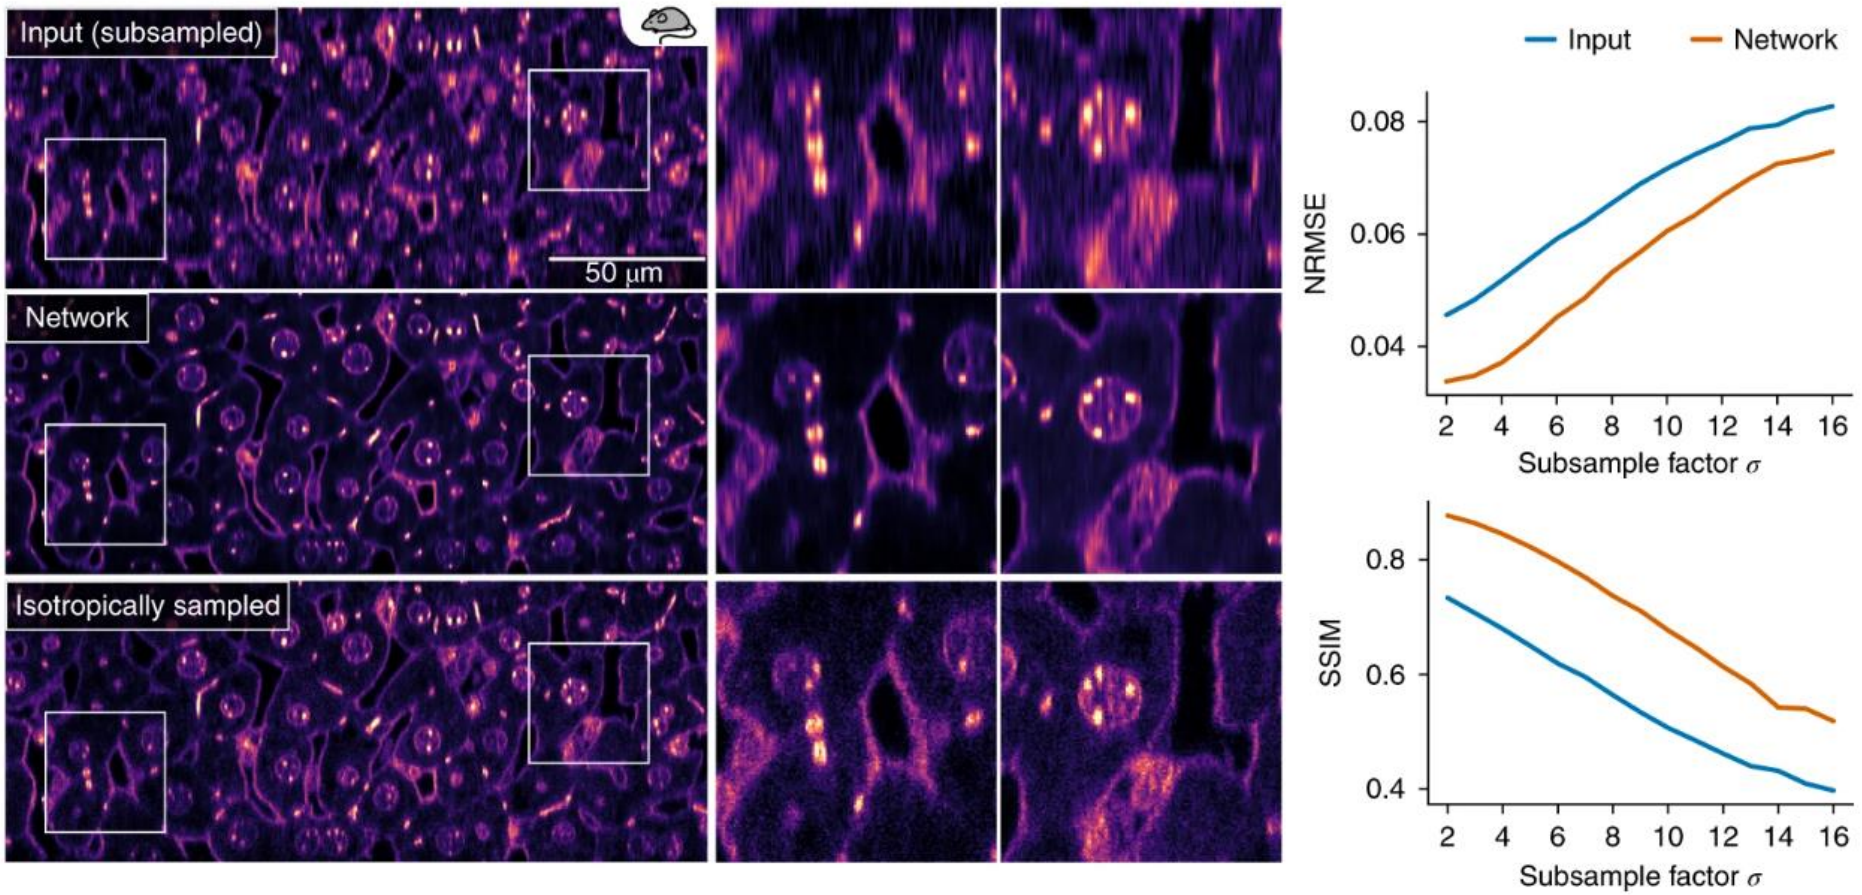

Weigert et al. (2018)

Anisotropic-to-isotropic restoration leads to improved segmentation

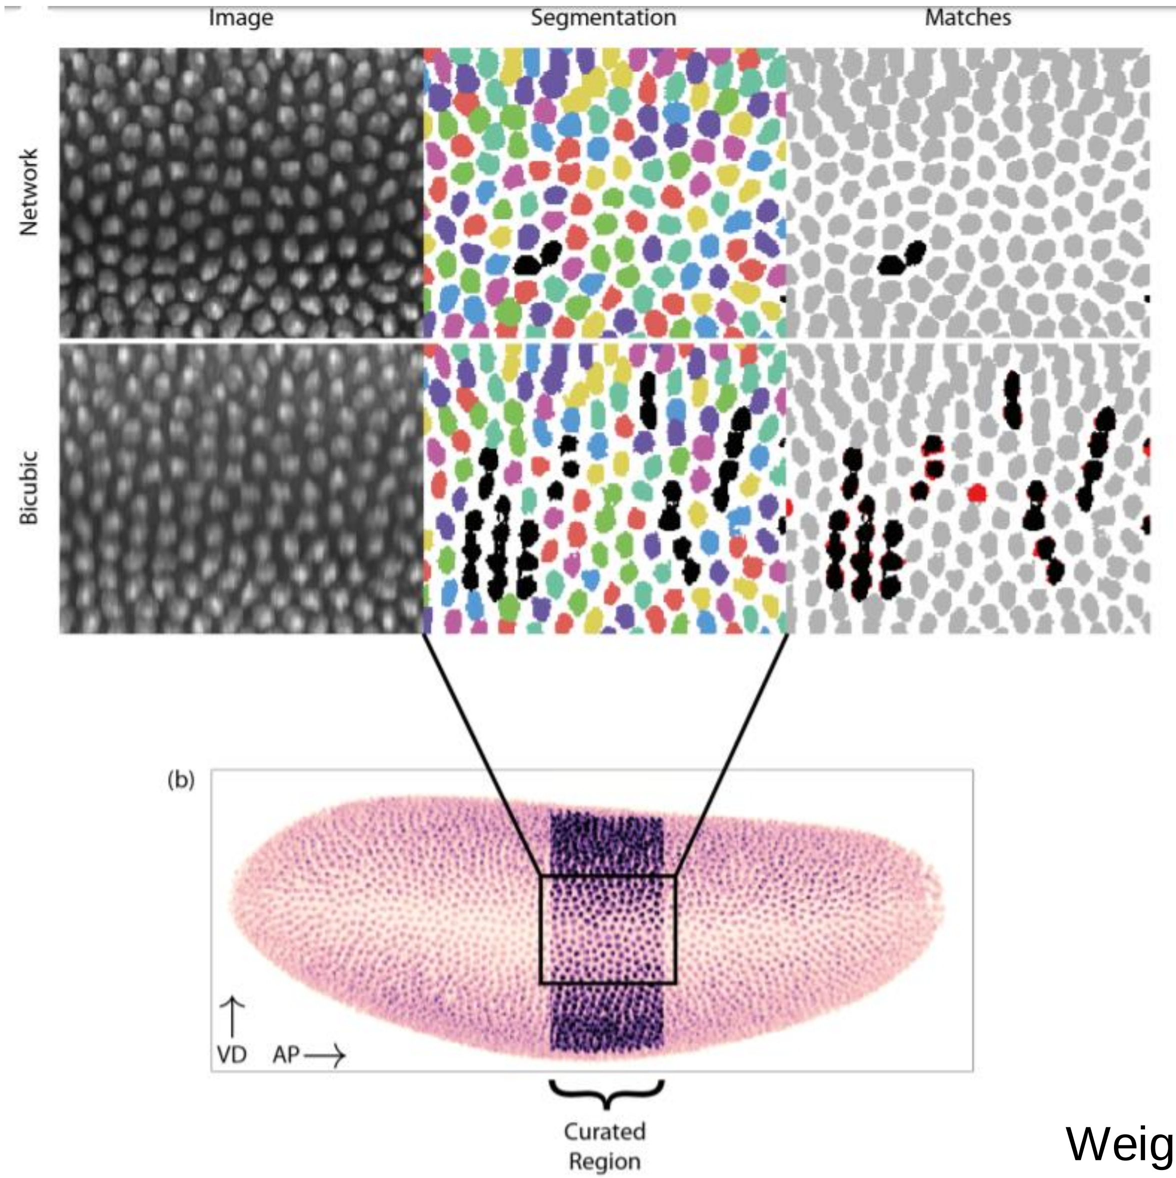

Weigert et al. (2018)

Anisotropic-to-isotropic restoration leads to improved segmentation

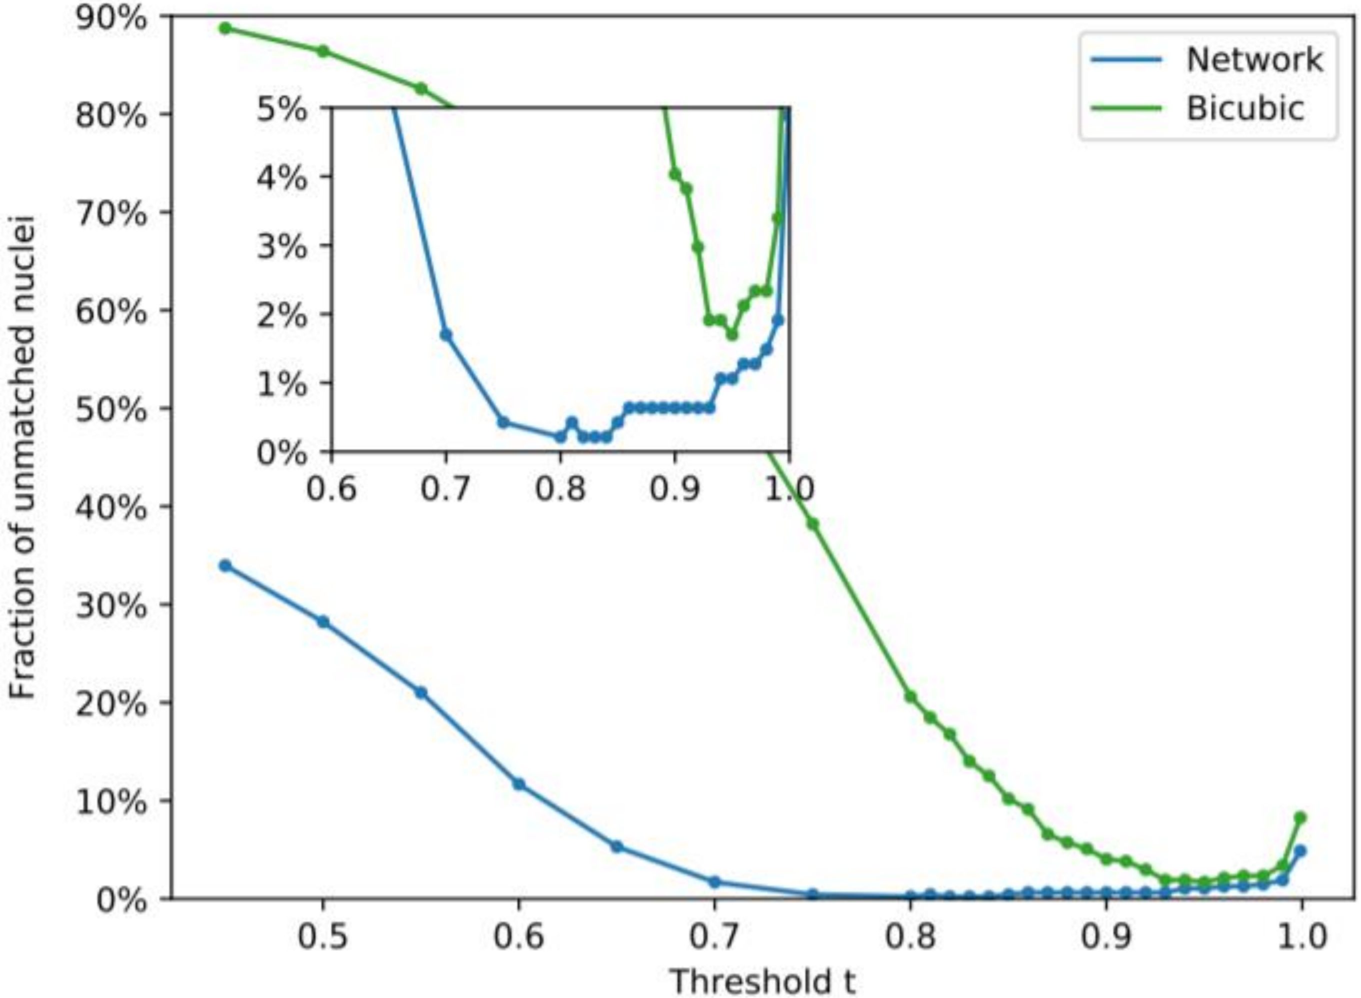

Weigert et al. (2018)

Synthetically generated training data

PSF, camera noise, and background auto-fluorescence

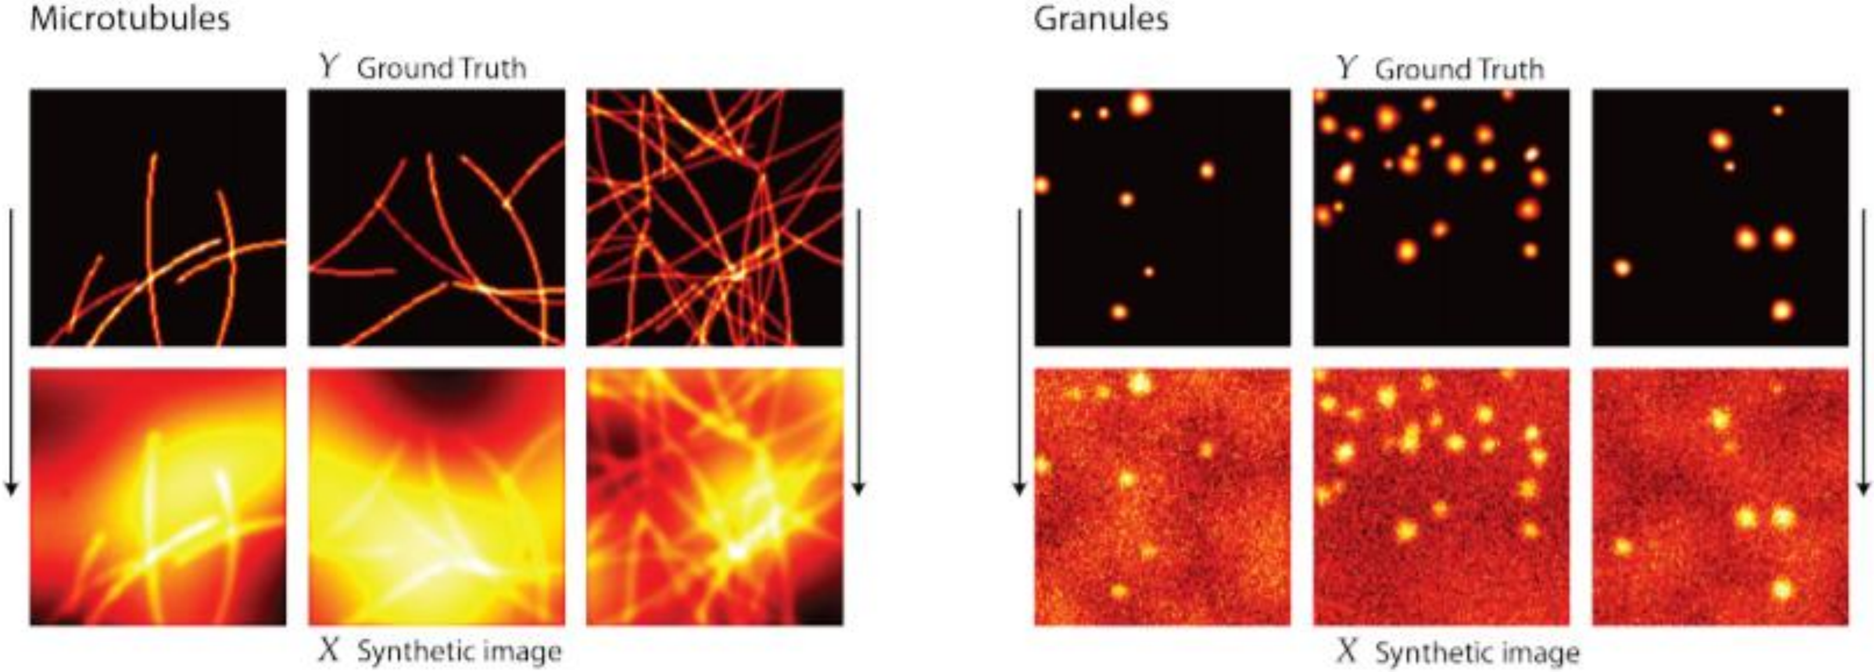

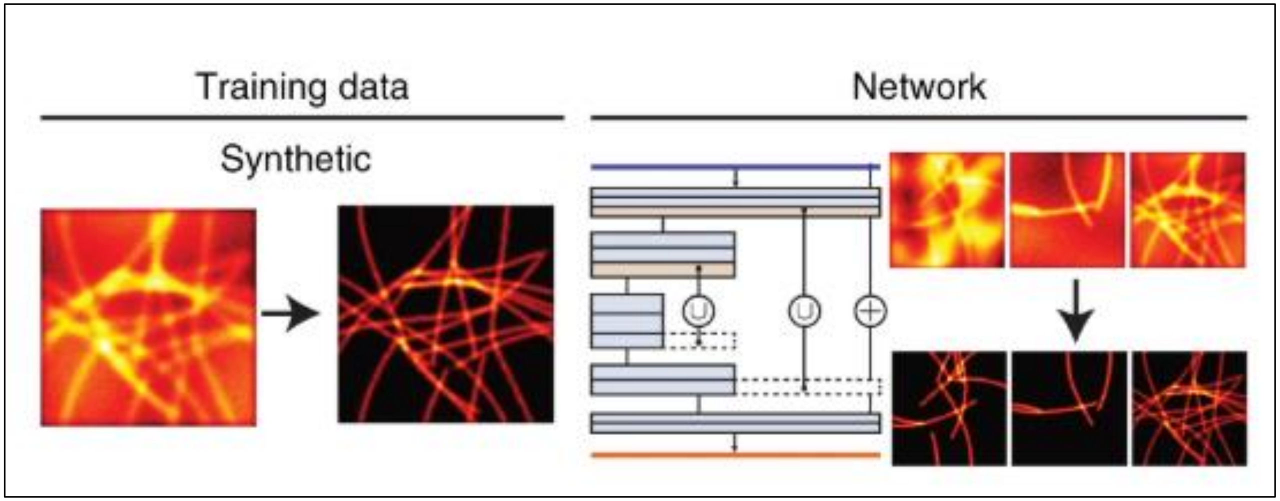

Weigert et al. (2018)

Image restoration with synthetic training data

Restoring sub-diffraction structures using only widefield images

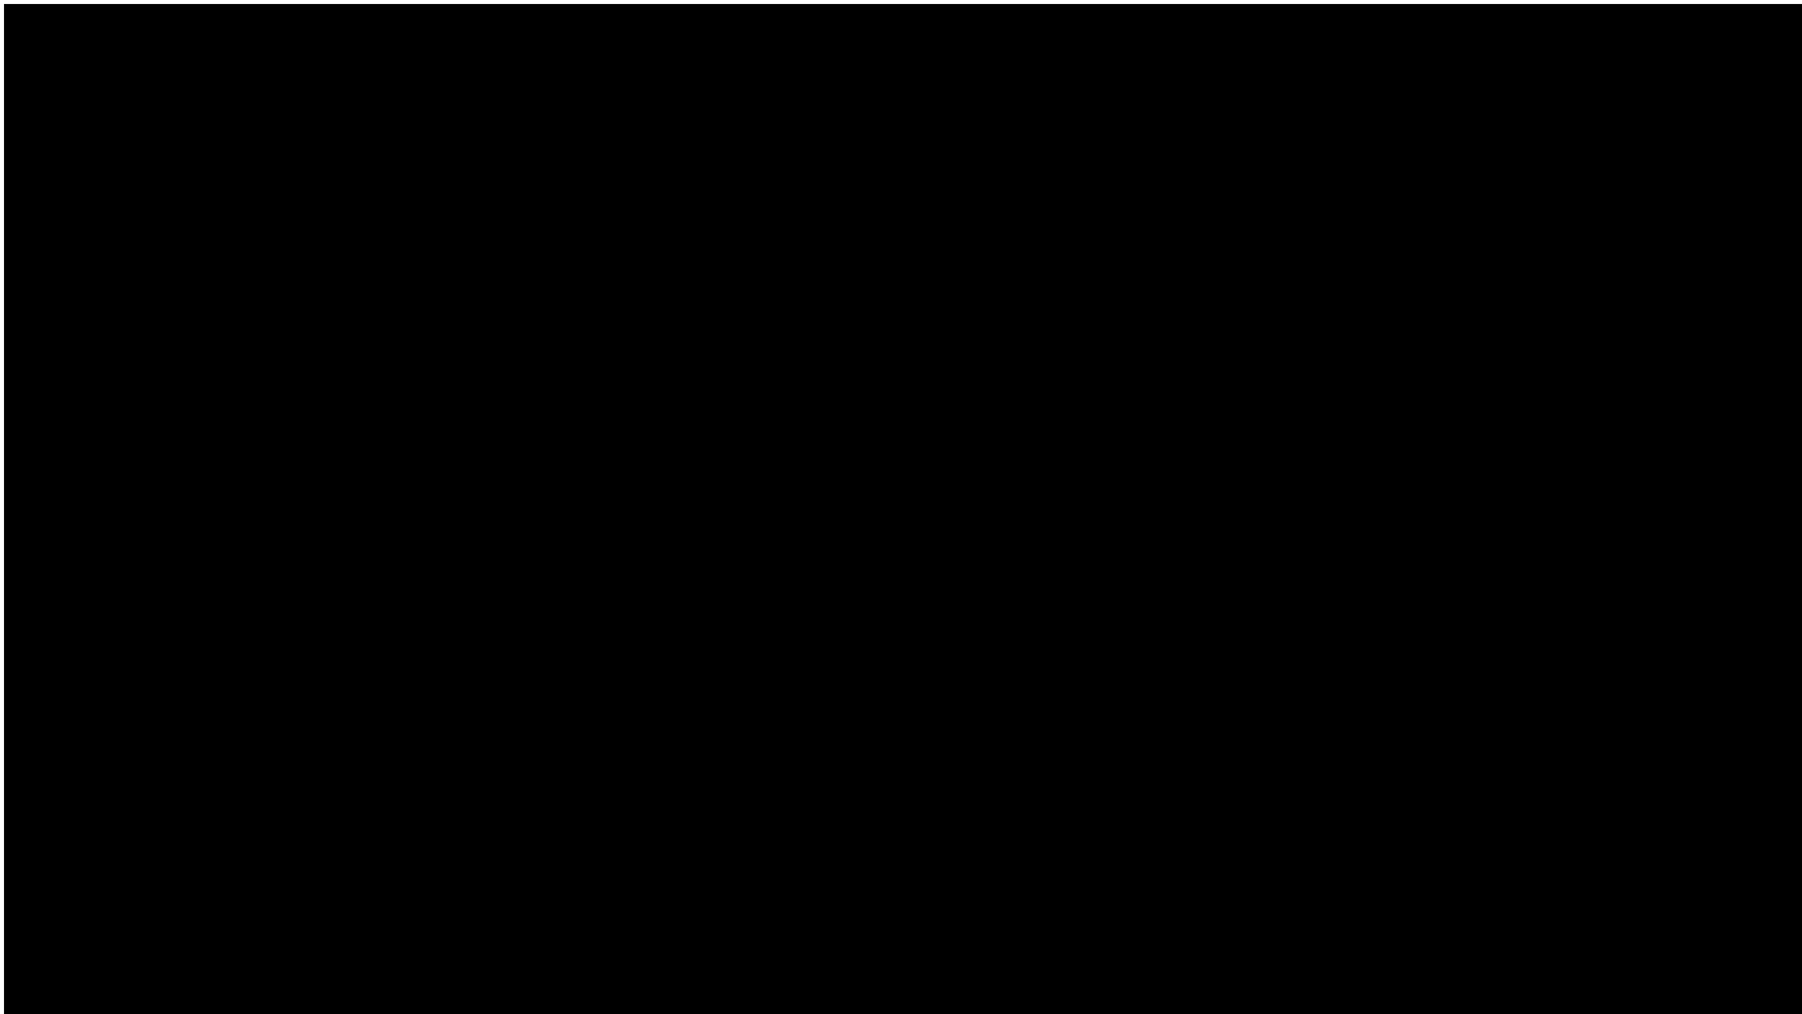

Weigert et al. (2018)

Improved performance of the CARE network over deconvolution

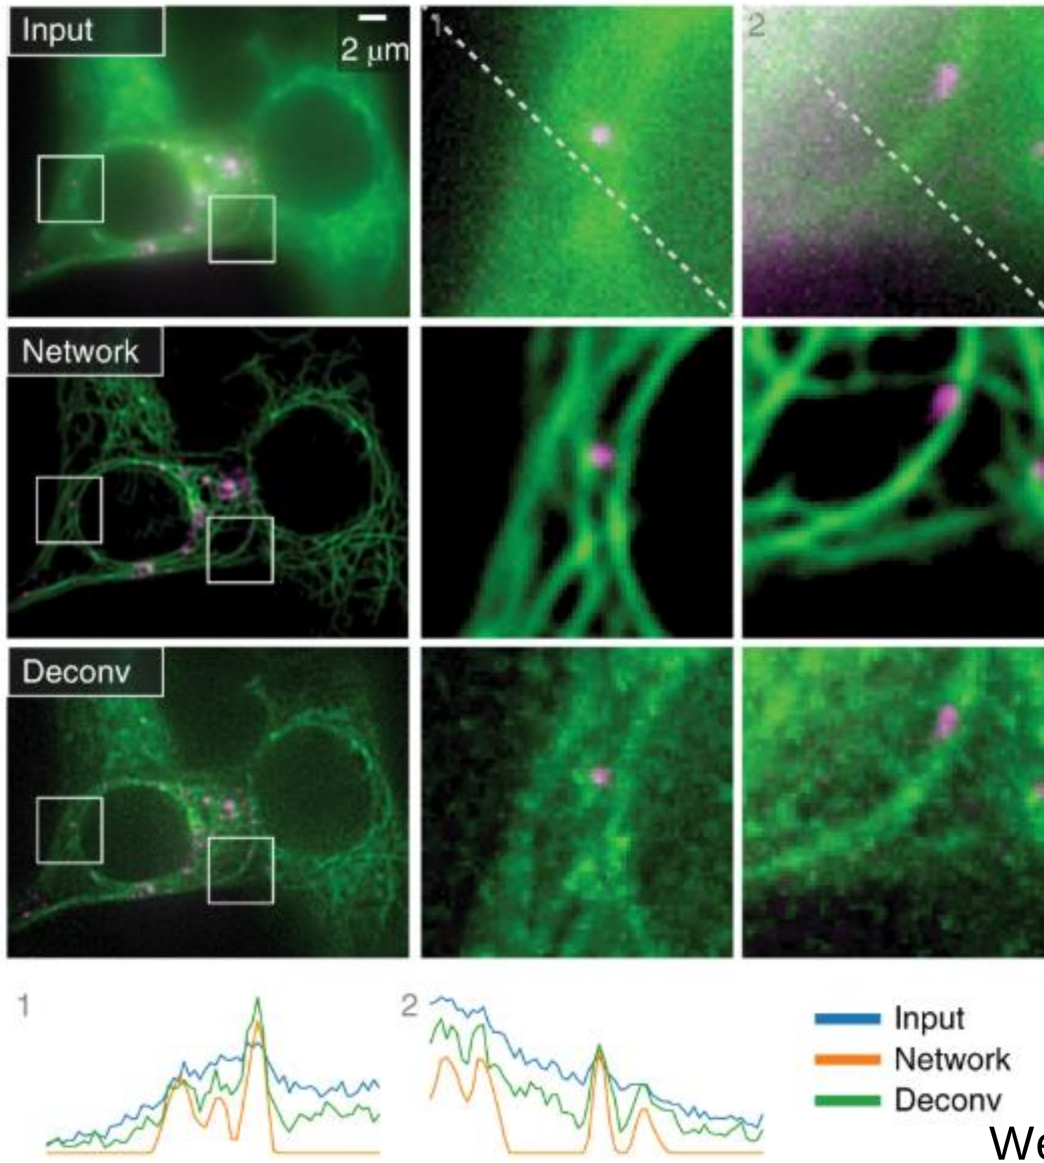

Weigert et al. (2018)

CARE vs. super-resolution radial fluctuations (SRRF)

Enabling 20-fold faster imaging

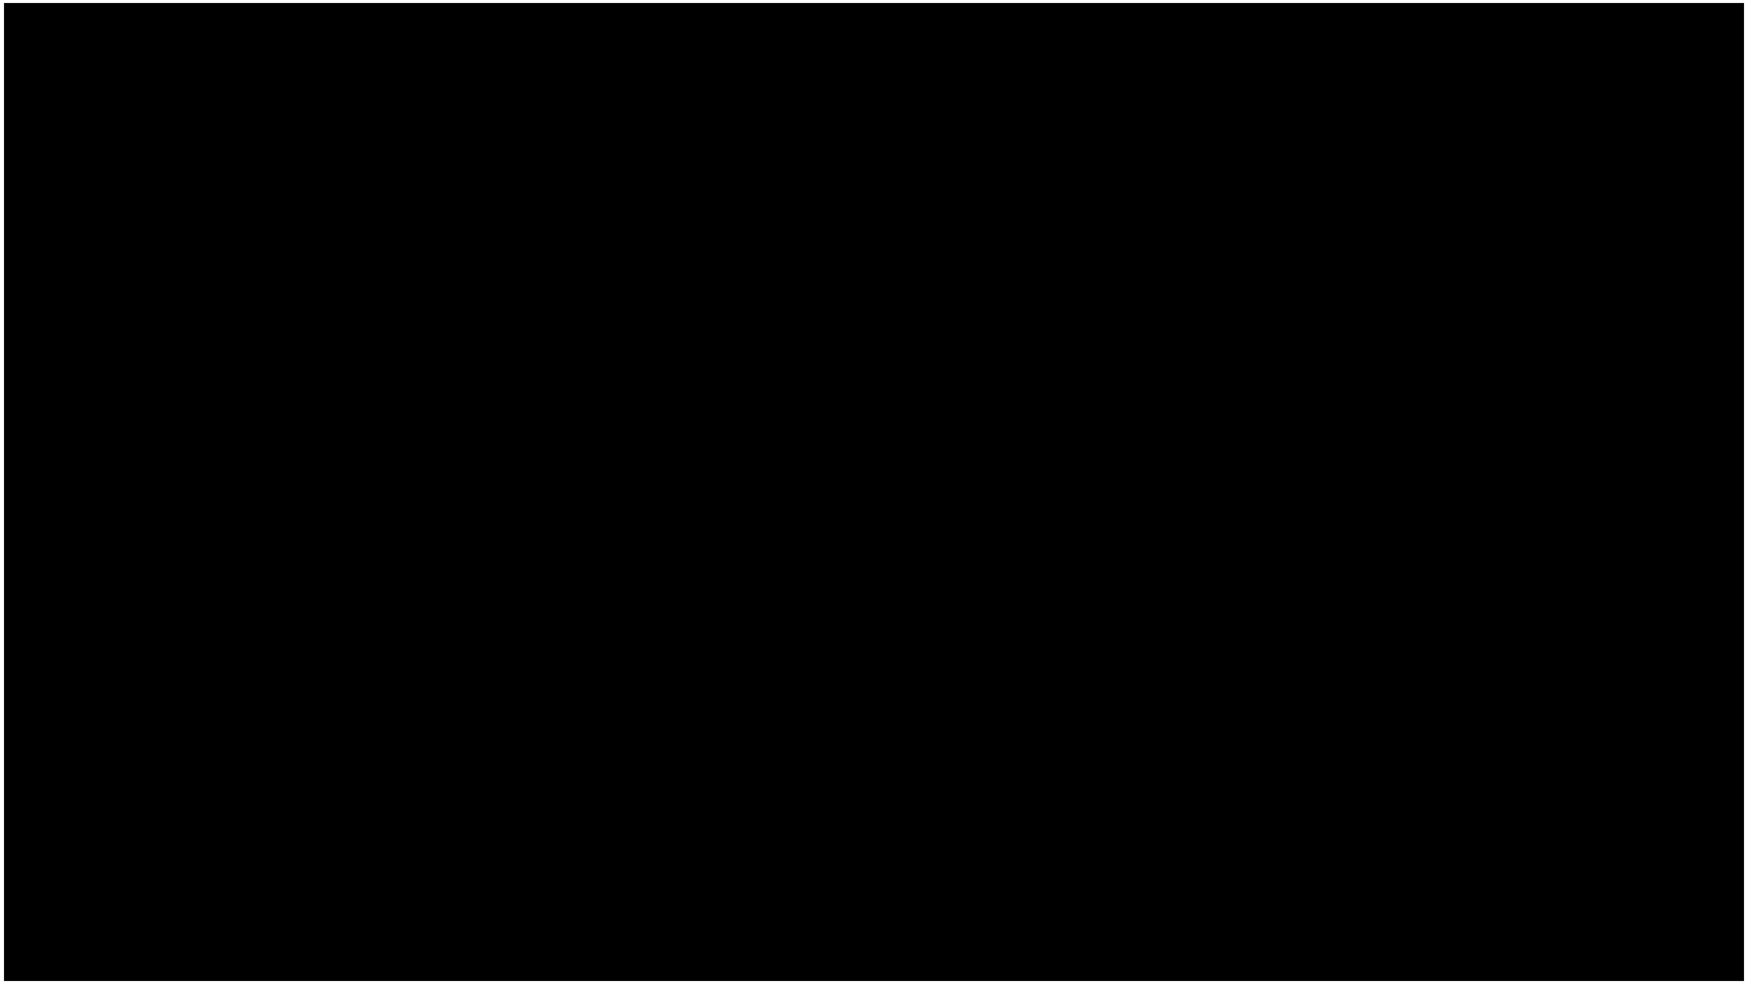

Weigert et al. (2018)

Cross-application of trained care networks

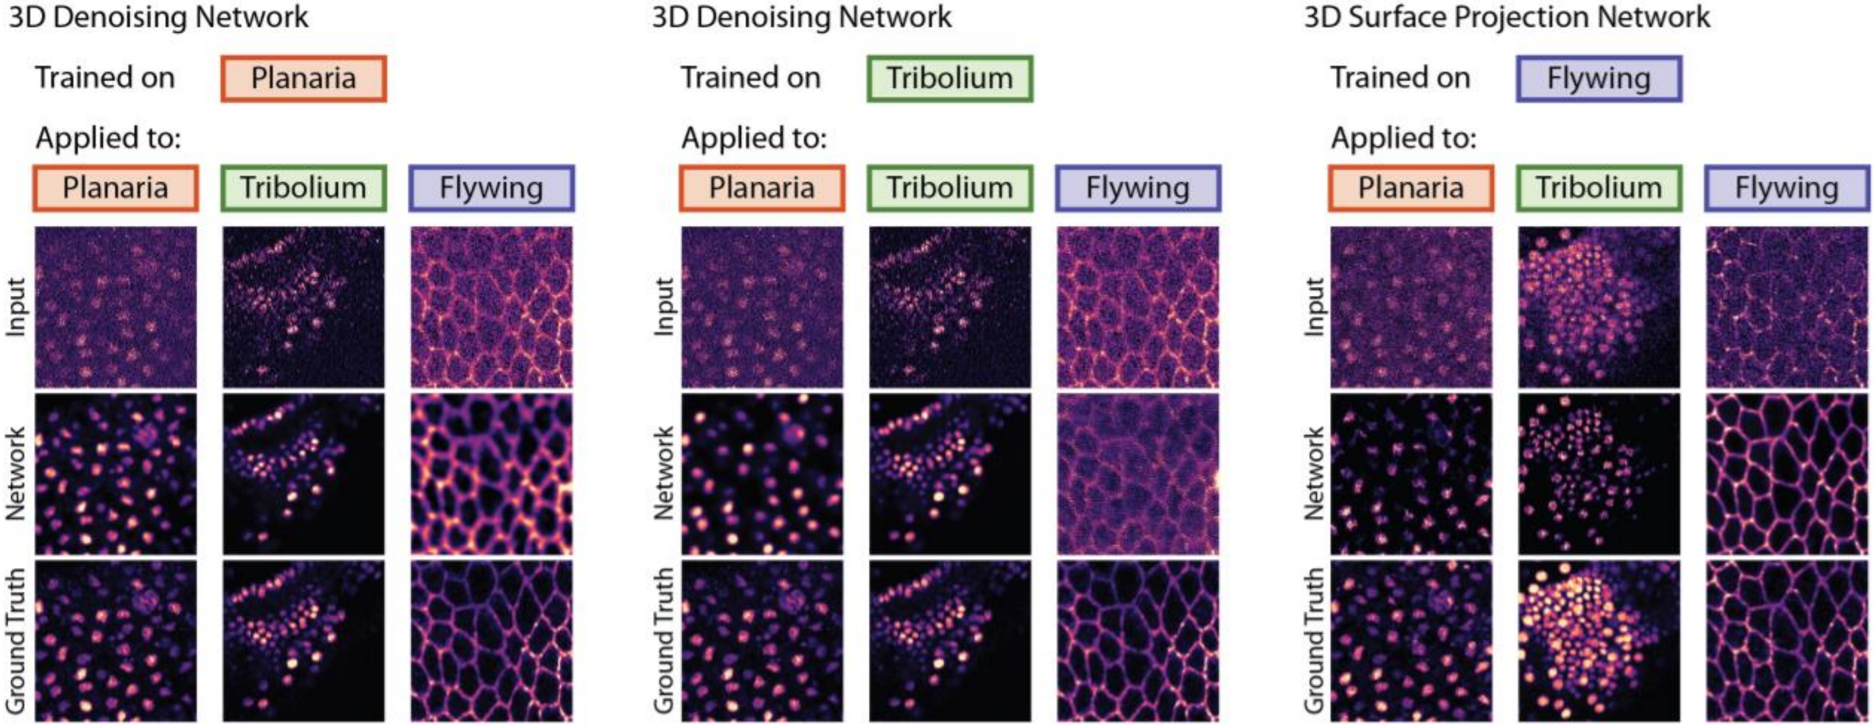

Weigert et al. (2018)

Minimal 'hallucination' effects

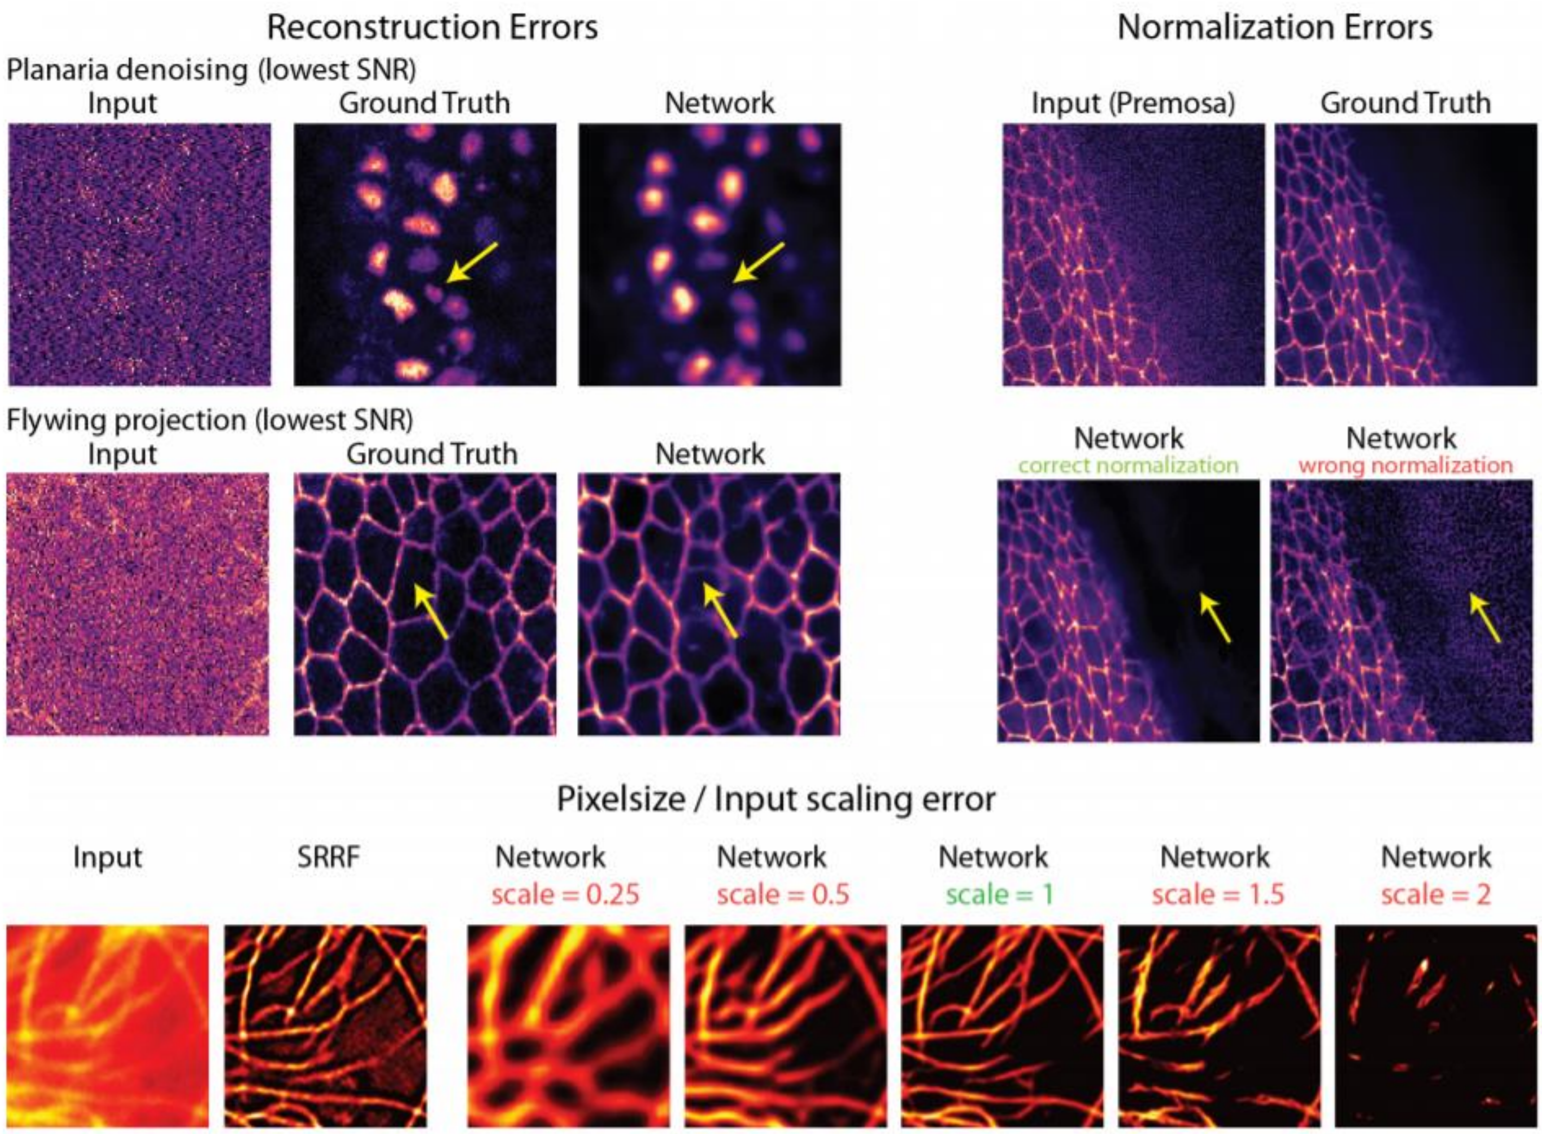

Weigert et al. (2018)

Reliability of image restoration

Calculating pixel-wise confidence intervals

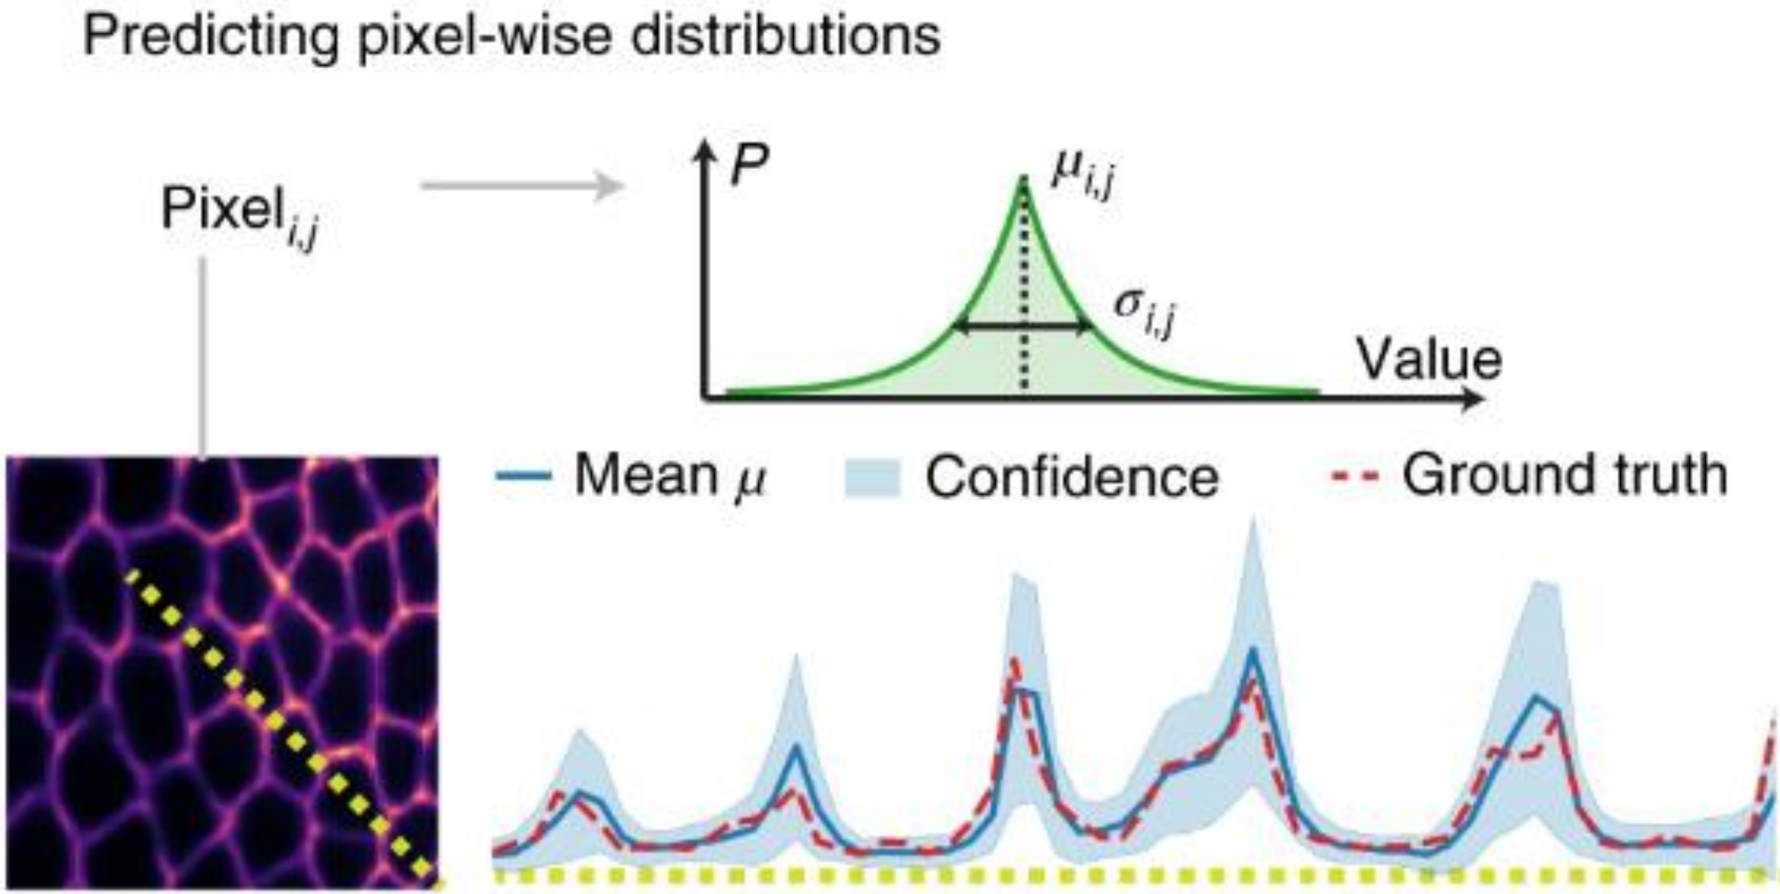

Weigert et al. (2018)

Visualizing the uncertainty by random sampling pixel intensities from their respective distributions

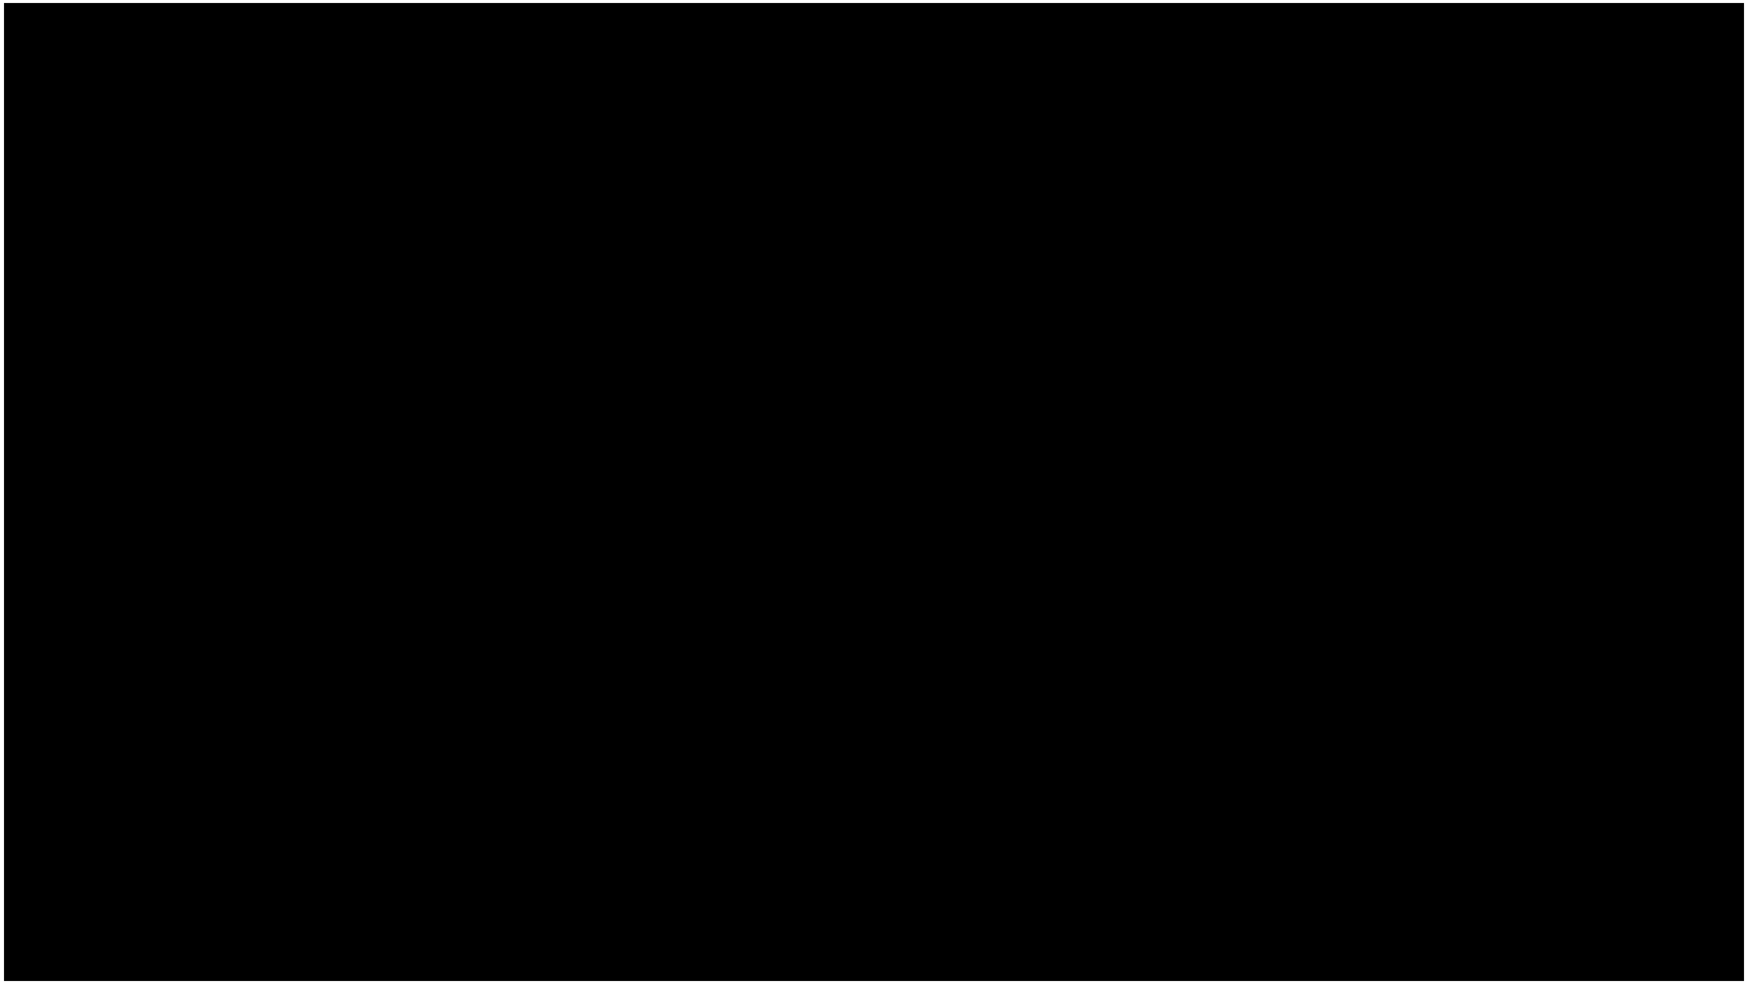

Weigert et al. (2018)

Assessing reliability via consistency of several trained models

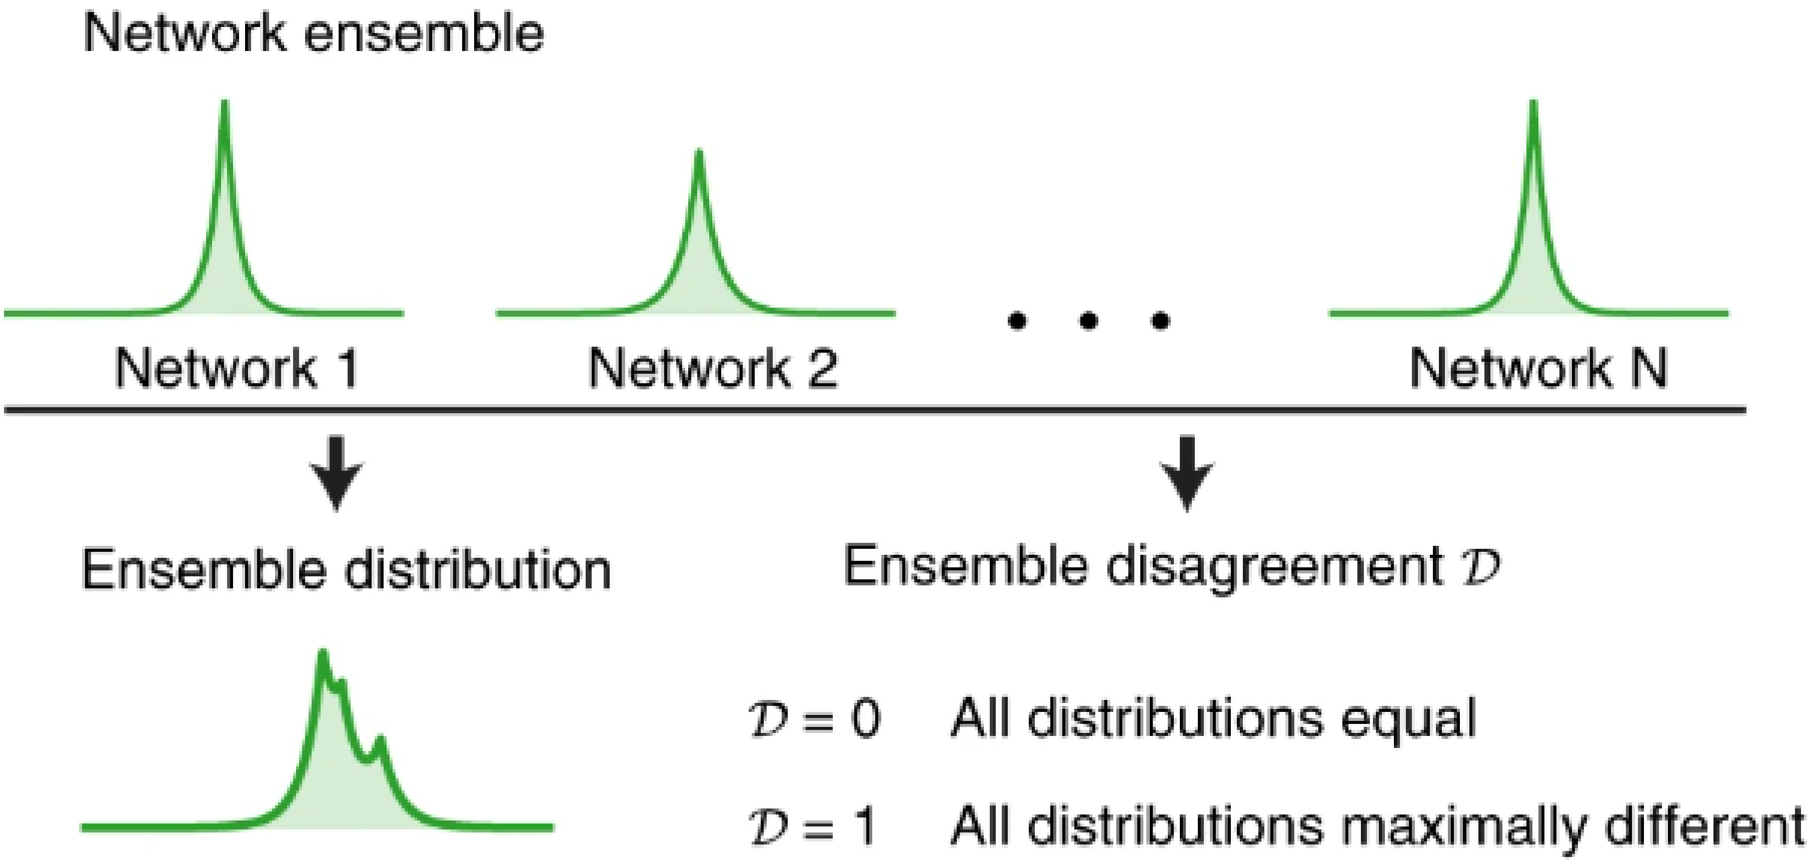

Weigert et al. (2018)

Assessing reliability via consistency of several trained models

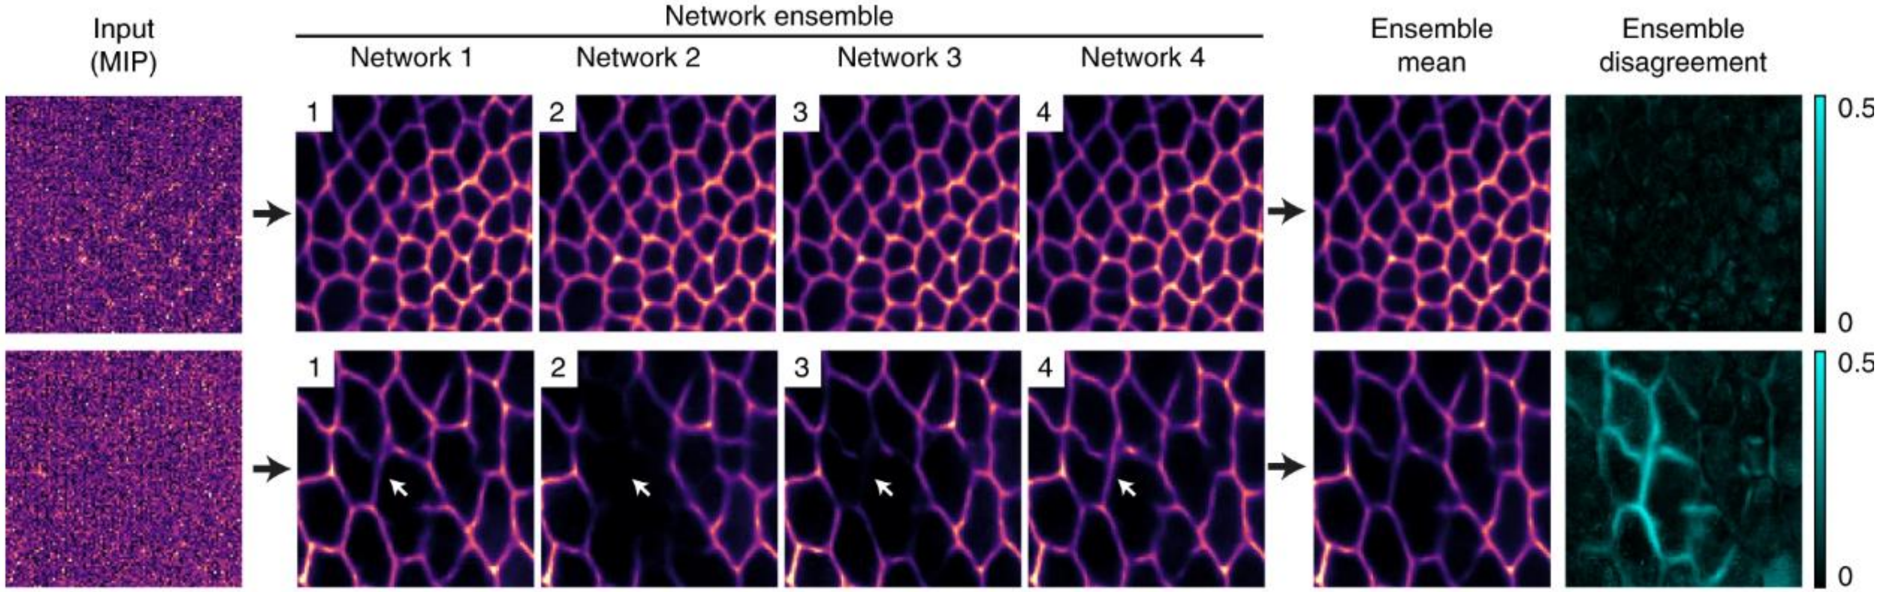

Weigert et al. (2018)

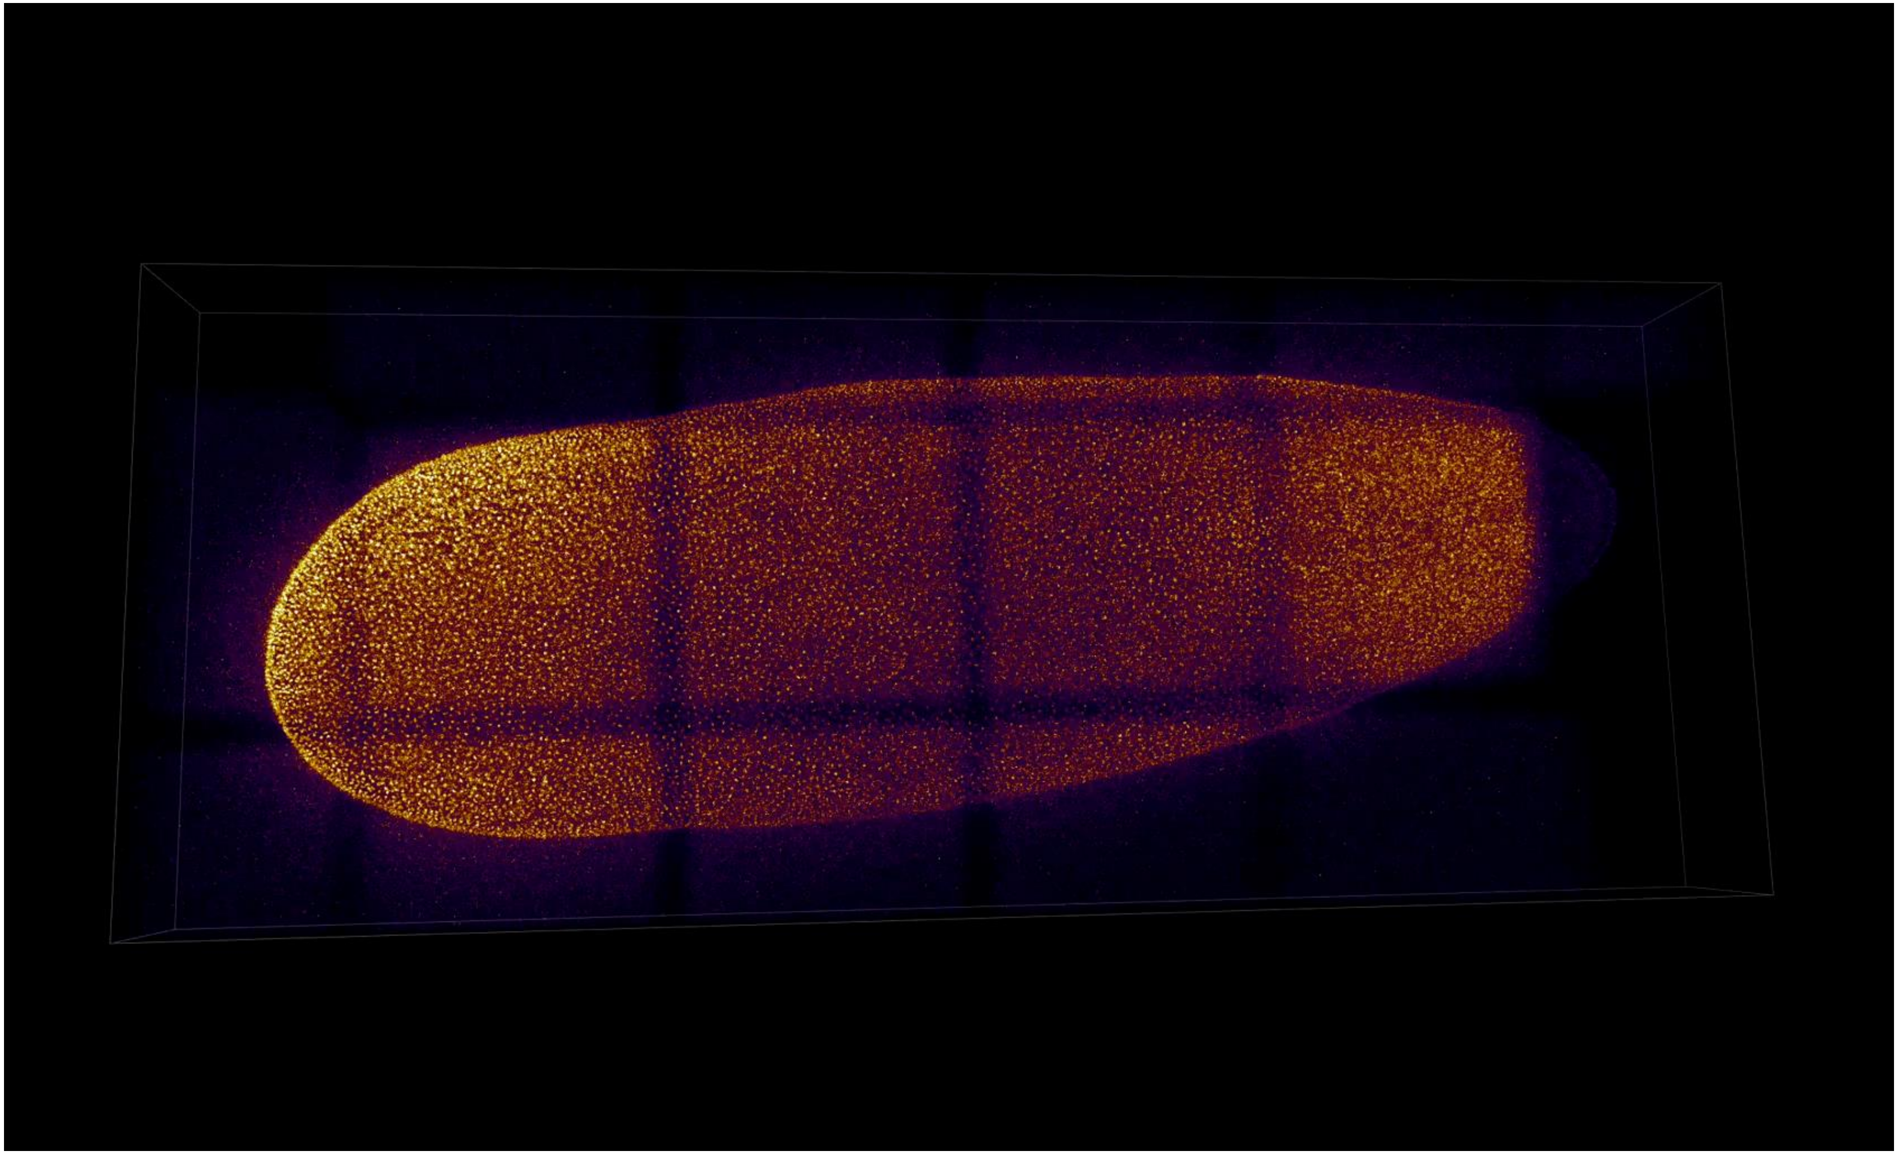

Weigert et al. (2018)

Usability is important!

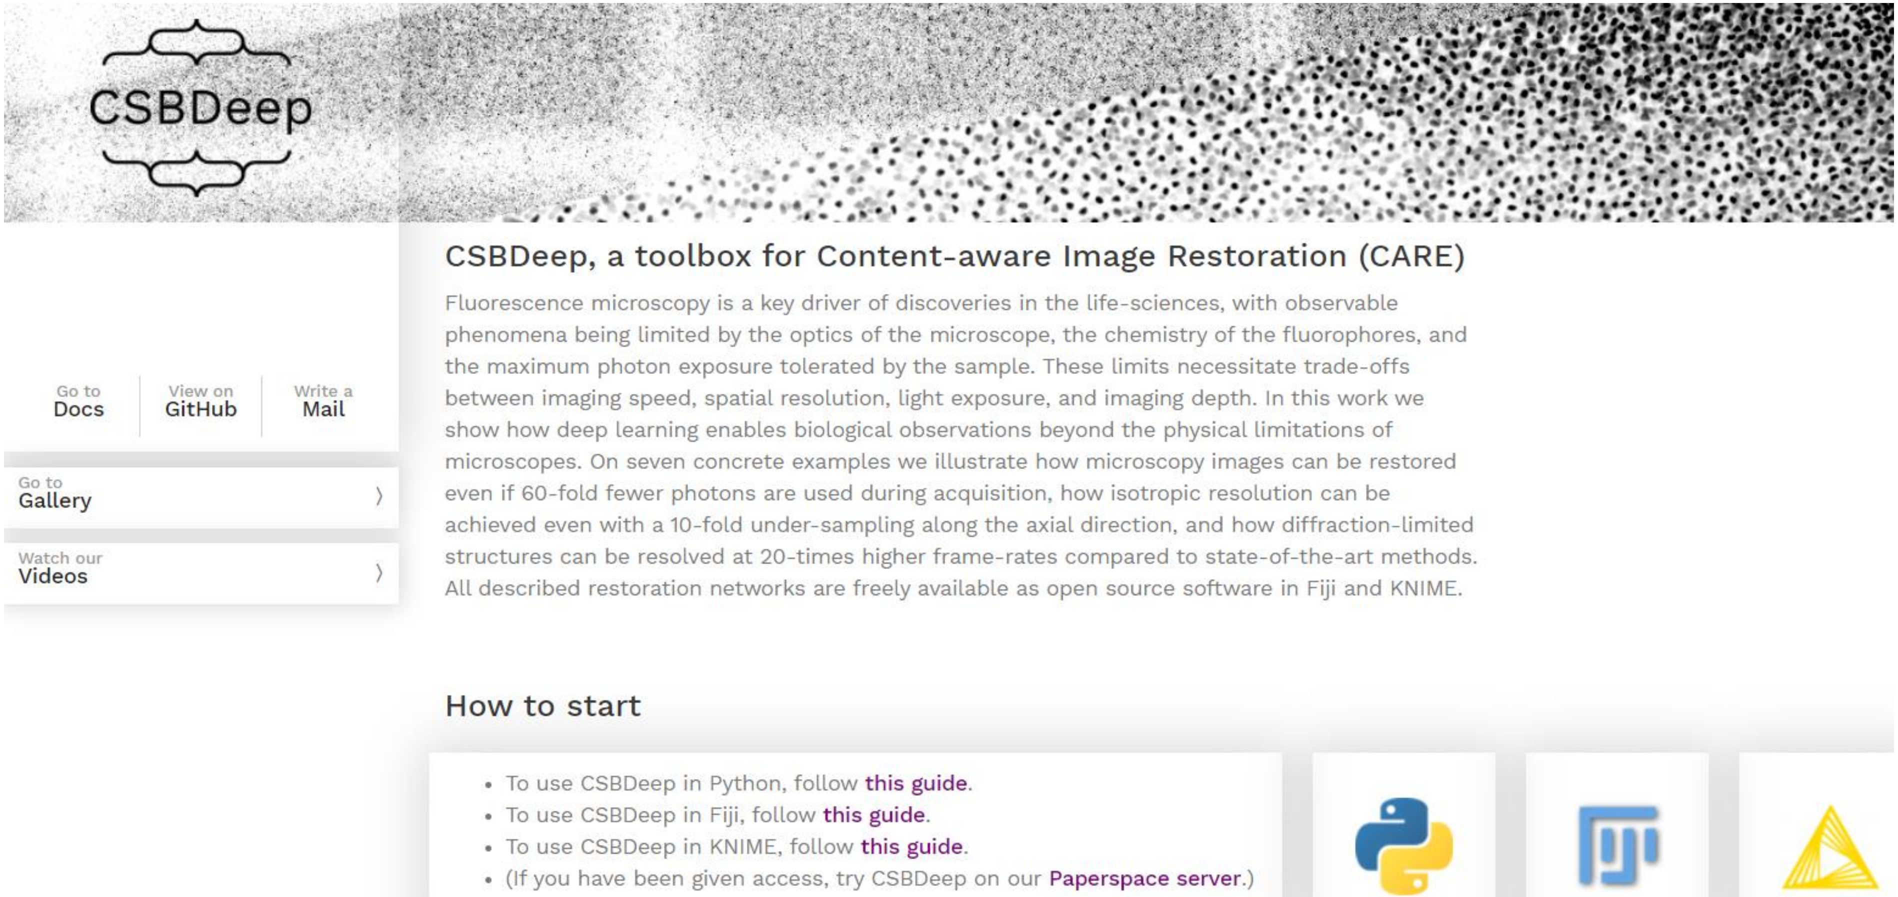

https://csbdeep.bioimagecomputing.com/
https://github.com/csbdeep/csbdeep

(Sometimes) even works out of the box!

Guillaume Jacquemet
@guijacquemet

#CARE vs #SRRF on my microtubule images (TIRF 100x). #CARE work beautifully. I look forward to train it to work on filopodia ! I already have images ready !
biorxiv.org/content/early/...

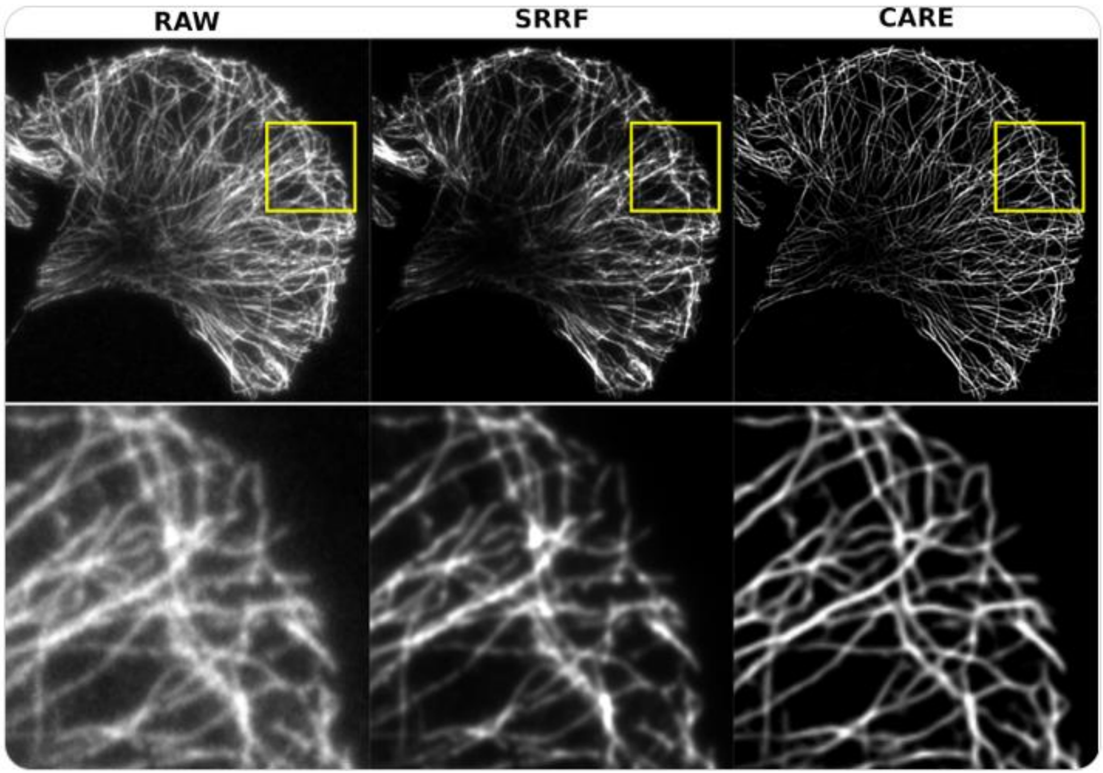

12:09 PM · Jan 4, 2018 · Twitter Web Client

Guillaume Jacquemet, https://twitter.com/guijacquemet/status/948859038795780101

Summary

- CARE leverages experiment-specific information to imaging tradeoffs
- Matched images, semi-synthetic and synthetic training data
- High quality denoising models can be trained **without the availability of clean ground truth data**
- Great for downstream analysis **not for intensity-based measurements!**
- Could work well "out of the box"
- Usable by others!

“Super resolution” in the context of computer vision

- Recover a high resolution image from one or more low resolution input images
- Classical multi-image super resolution: combining images obtained at subpixel misalignments (Irani & Peleg, 1991)
- Example-Based super-resolution: learning correspondence between matching low and high resolution image patches (Freeman, Pasztor & Carmichael, 2000)

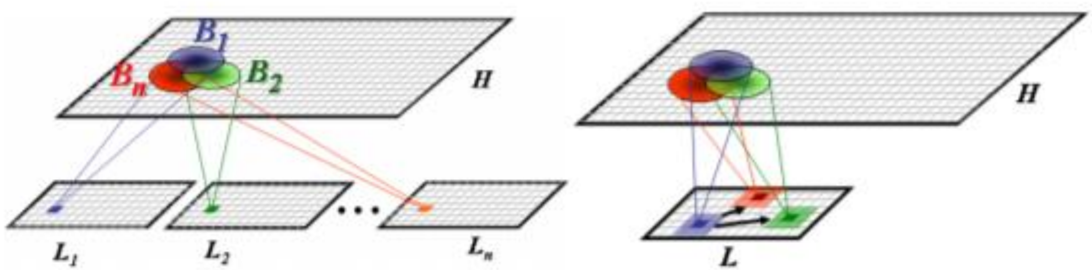

(a) Classical Multi-Image SR (b) Single-Image Multi-Patch SR

Glasner, Bagon and Irani (2009)

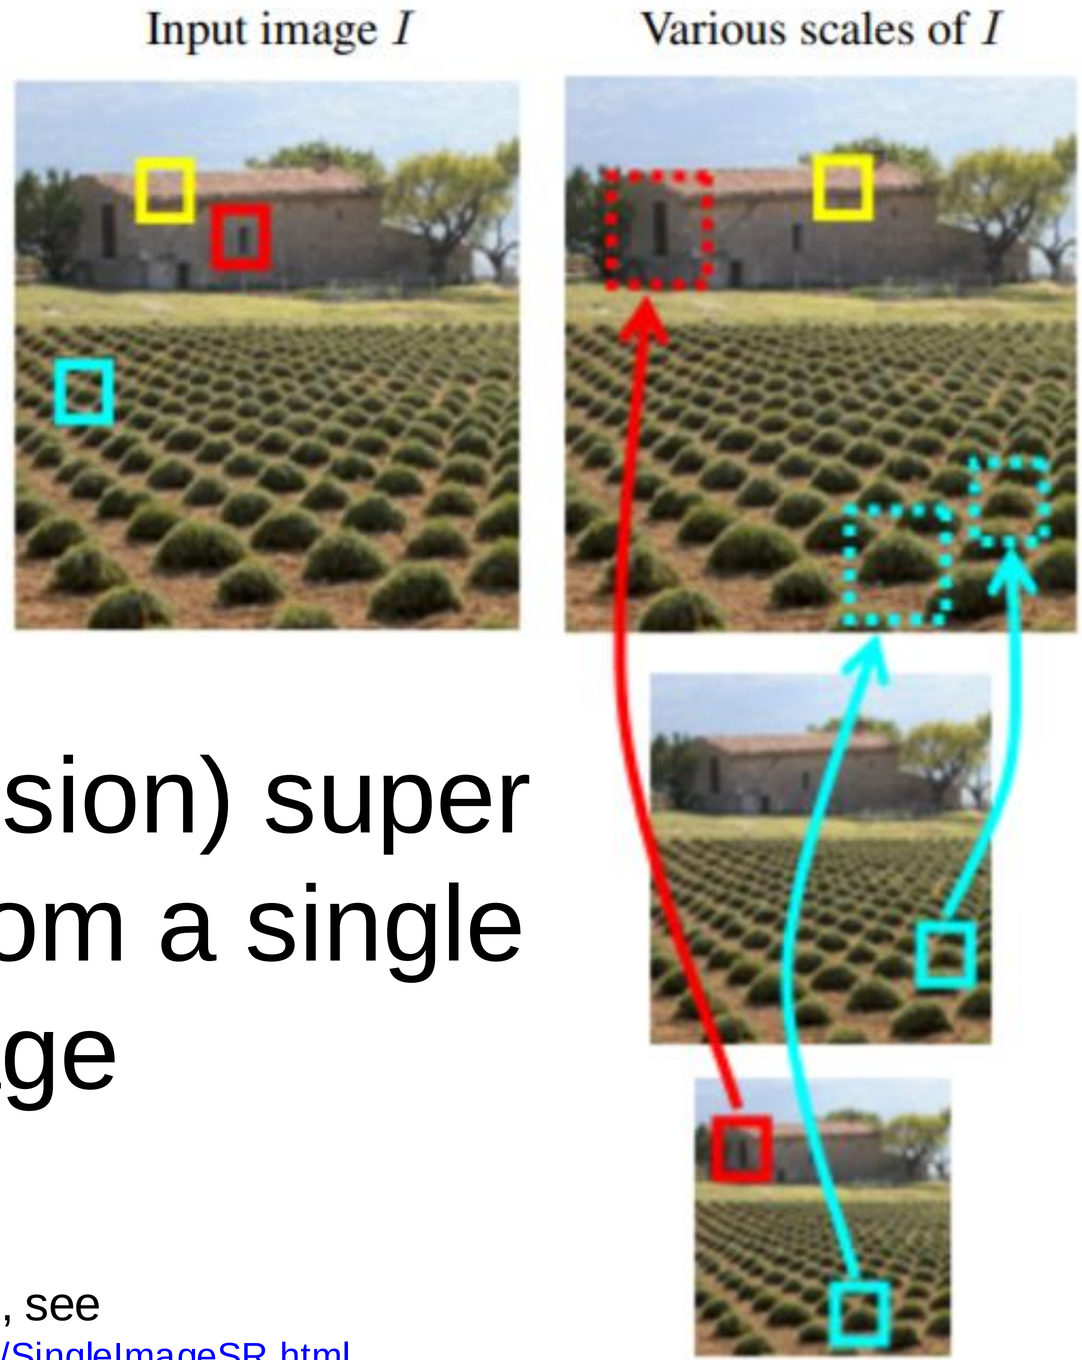

(computer vision) super resolution from a single image

Glasner, Bagon and Irani (2009), see
http://www.wisdom.weizmann.ac.il/~vision/SingleImageSR.html

(One) result

(a) Input image (scaled for display). (b) Bicubic interpolation (×2). (c) Within image repetitions (×2). (d) Unified single-image SR (×2).

Glasner, Bagon and Irani (2009), see http://www.wisdom.weizmann.ac.il/~vision/SingleImageSR.html

2009 vs. 2018

A conceptual lag of (almost) a **decade** between computer vision and microscopy!

There might be other opportunities out there..

Handling noisy images: new extensions from the Jug lab

(could be picked as a student presentation)

- Noise2noise: a mapping between pairs of independently degraded versions of the same training image, requires availability of pairs of noisy images (Lehtinen et. al., 2018)
- Noise2Void: <u>self-supervised</u> training method with "blind-spot" networks – excluding a pixel and learning a mapping from the noisy image to the missing pixel (Krull, Buchholz and Jug, 2019)
- Probabilistic Noise2Void: predict per-pixel intensity distributions (Krull, Viˇcar and Jug, 2019)

Examples of mapping matched image pairs for image enhancement

- Super resolution: <u>ANNA-PALM</u>, deep STORM, "cross modality" (resolution), PSSR
- 2D → 3D refocusing, Deep-Z

Could be picked as a student presentation

ANNA-PALM: DL accelerates super-resolution localization microscopy

- ANNA-PALM (slightly) preceded CARE
- exploits the structural redundancy of most biological images to reconstruct high-quality images from under-sampled localization microscopy data
- Reducing total number of frames and independent localizations without trading off spatial resolution

Ouyang et al. (2018)

Photoactivated localization microscopy (PALM)

diffraction-limited image sequences | Widefield | Localization

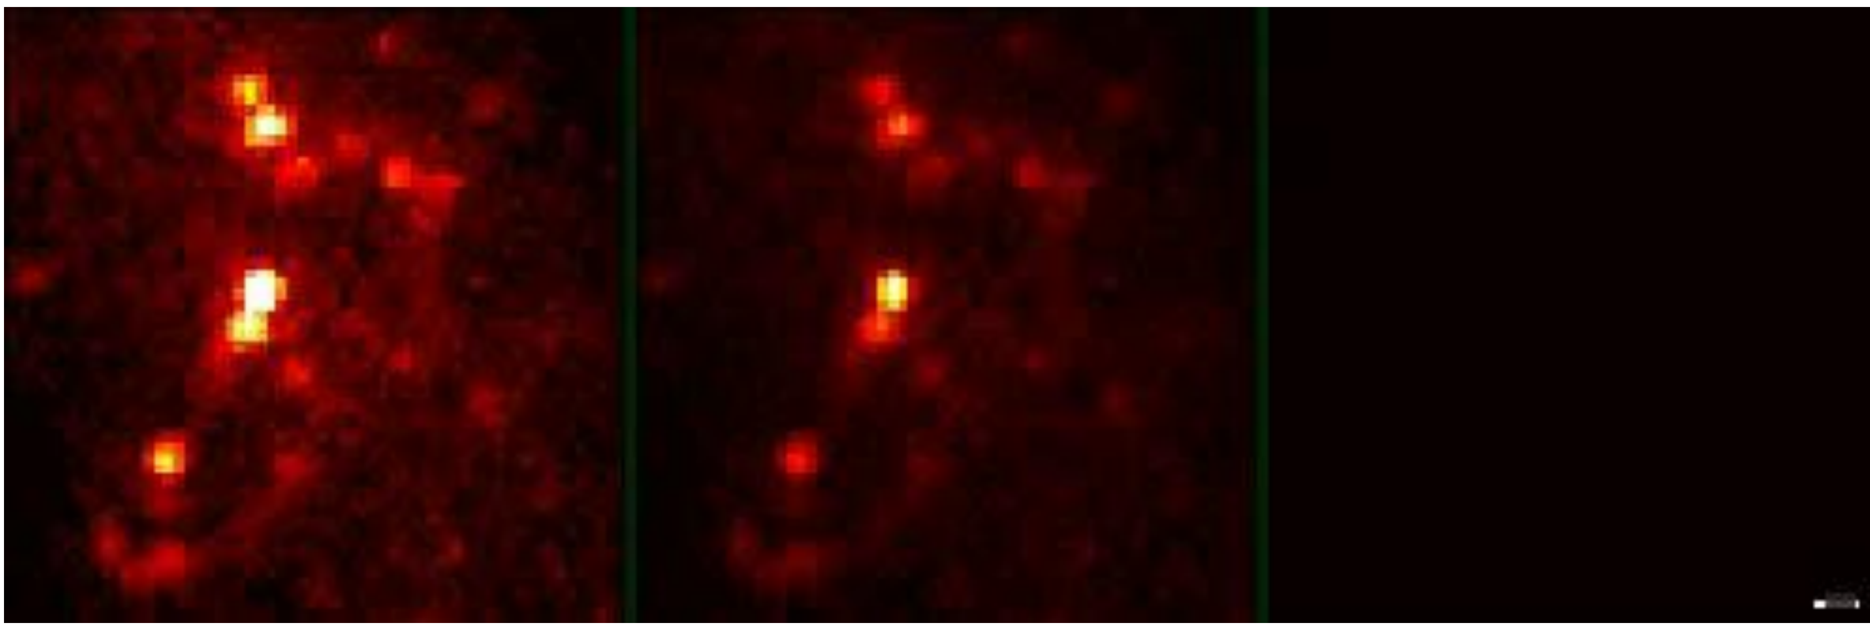

Betzig (2006)

ANNA-PALM network training

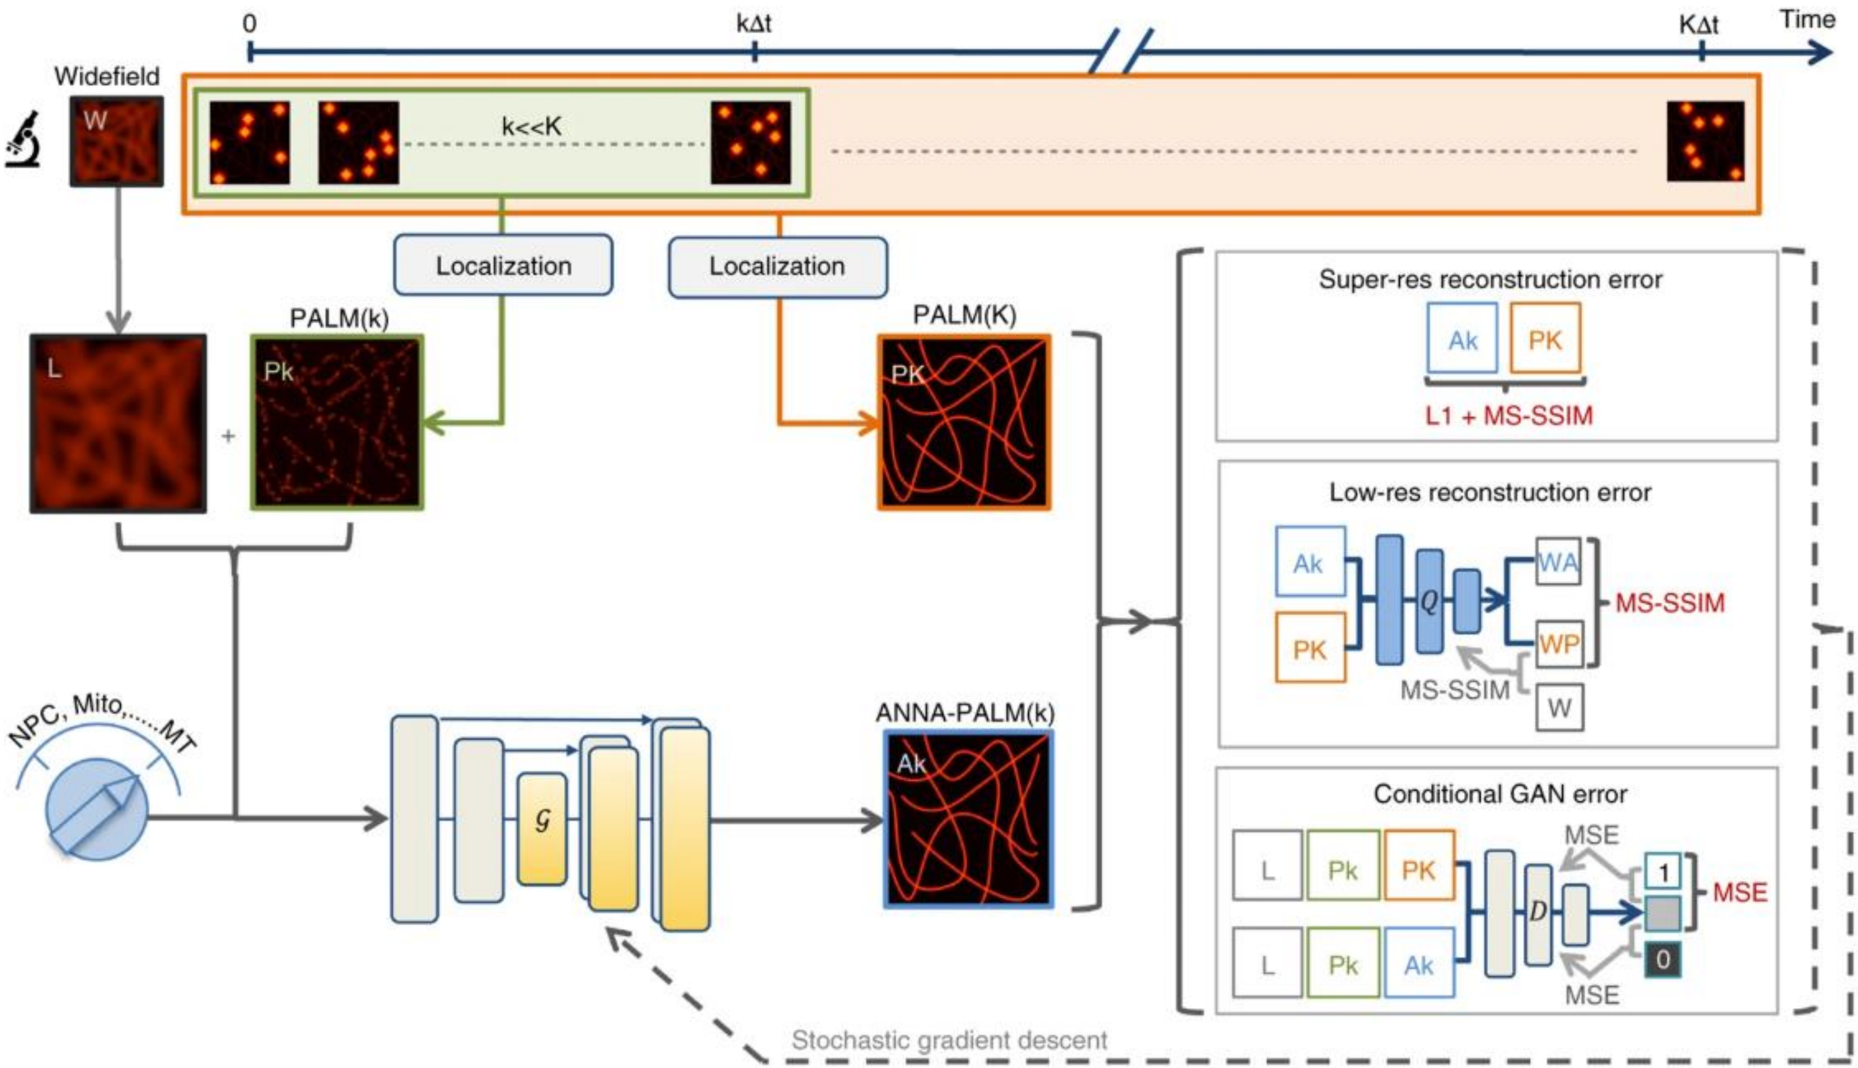

Ouyang et al. (2018)

ANNA-PALM inference

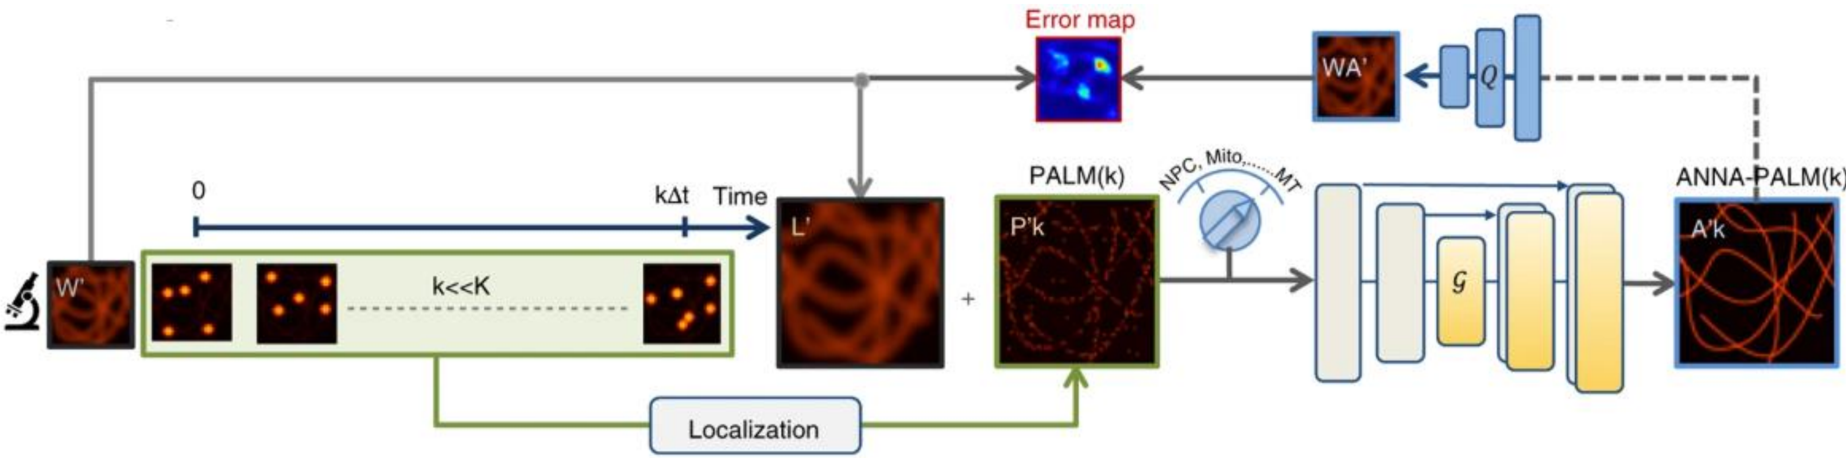

Ouyang et al. (2018)

Validation of ANNA-PALM on simulated images

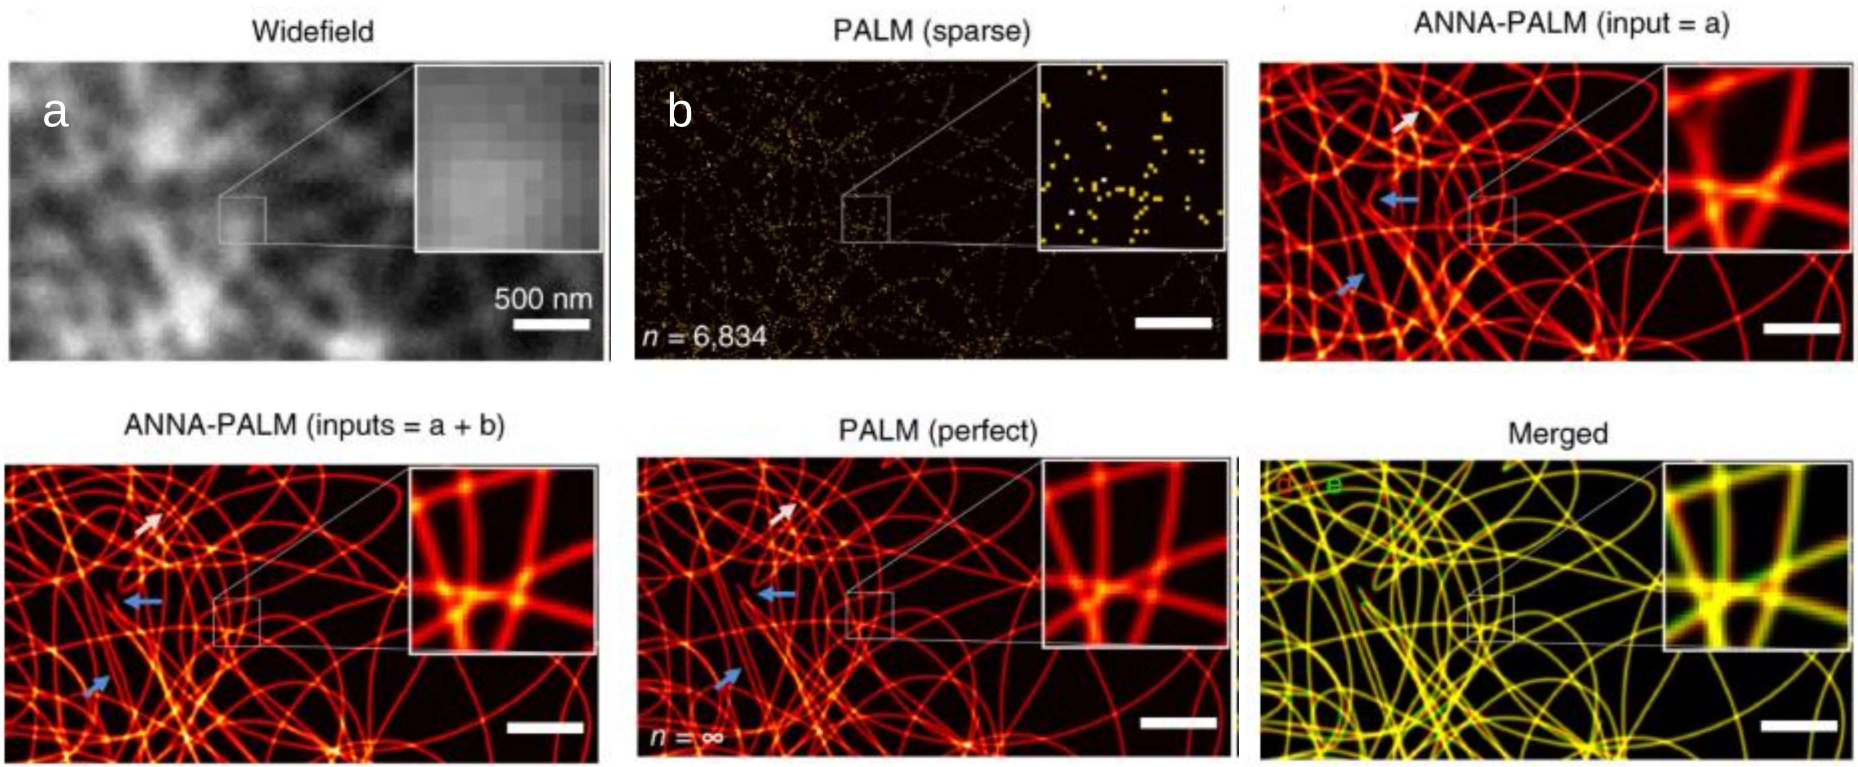

Ouyang et al. (2018)

Validation of ANNA-PALM on simulated images

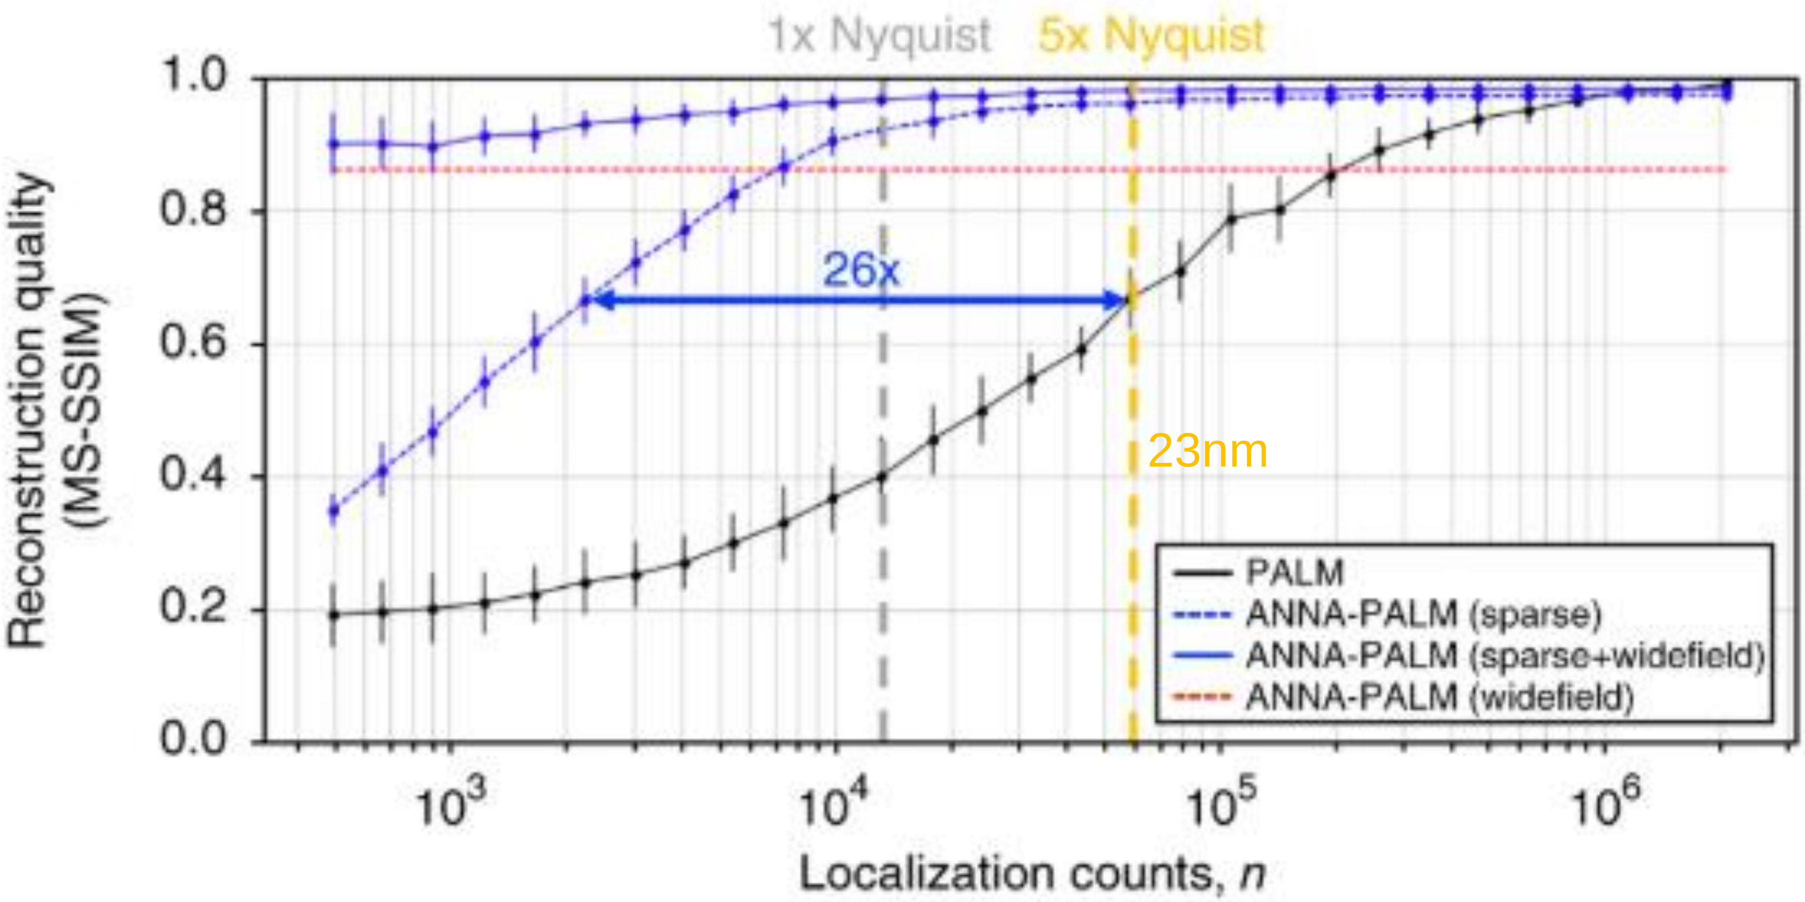

Ouyang et al. (2018)

ANNA-PALM imaging of microtubules

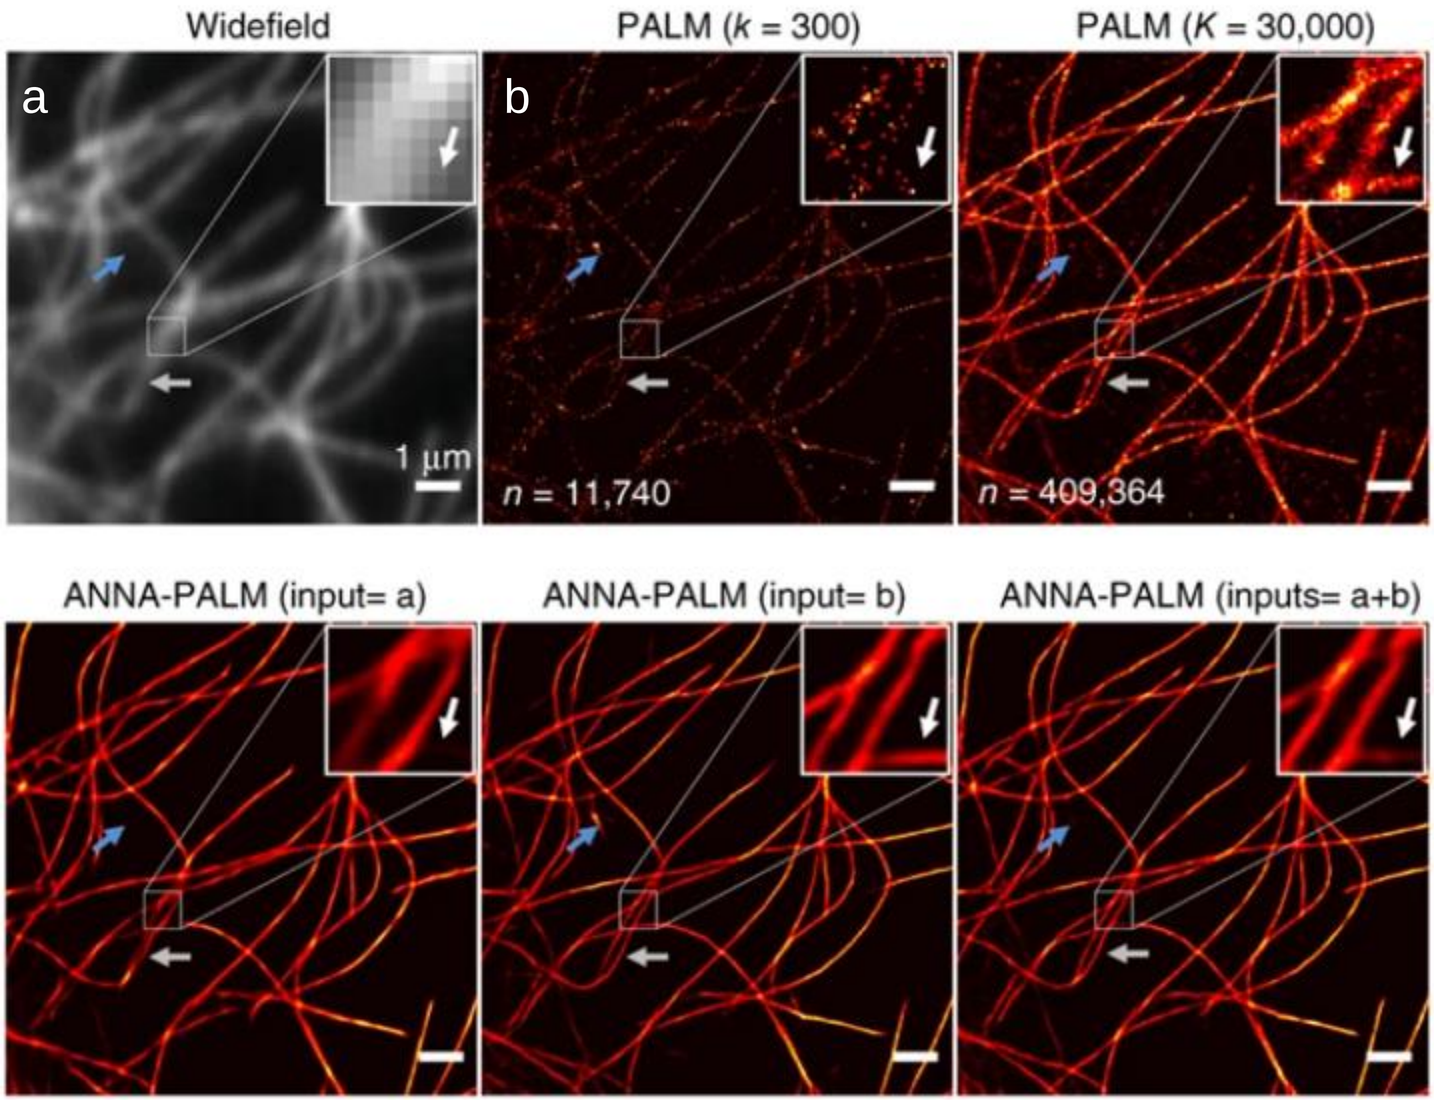

Ouyang et al. (2018)

ANNA-PALM imaging of microtubules

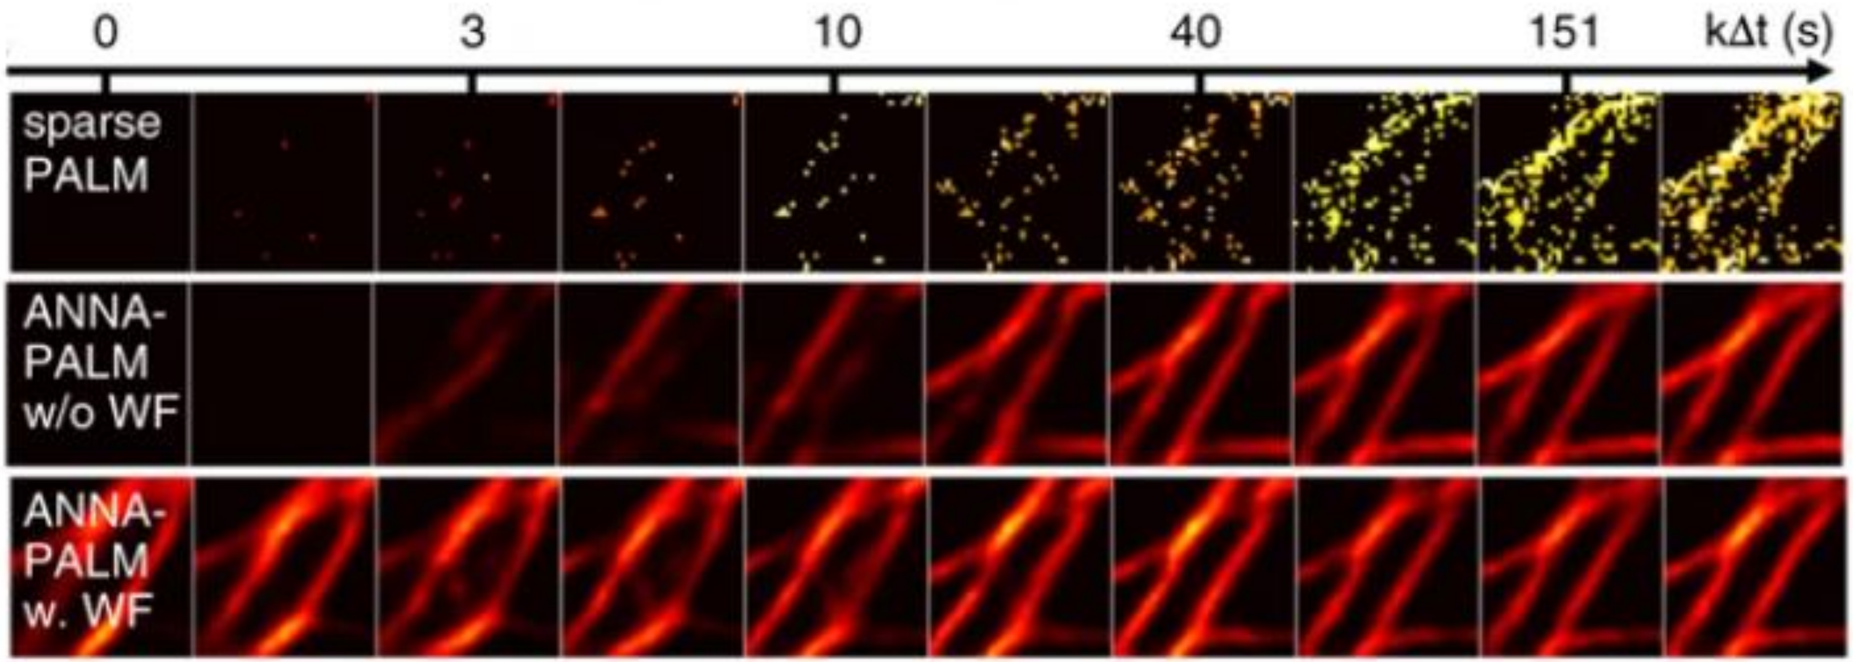

Ouyang et al. (2018)

High-throughput super-resolution imaging

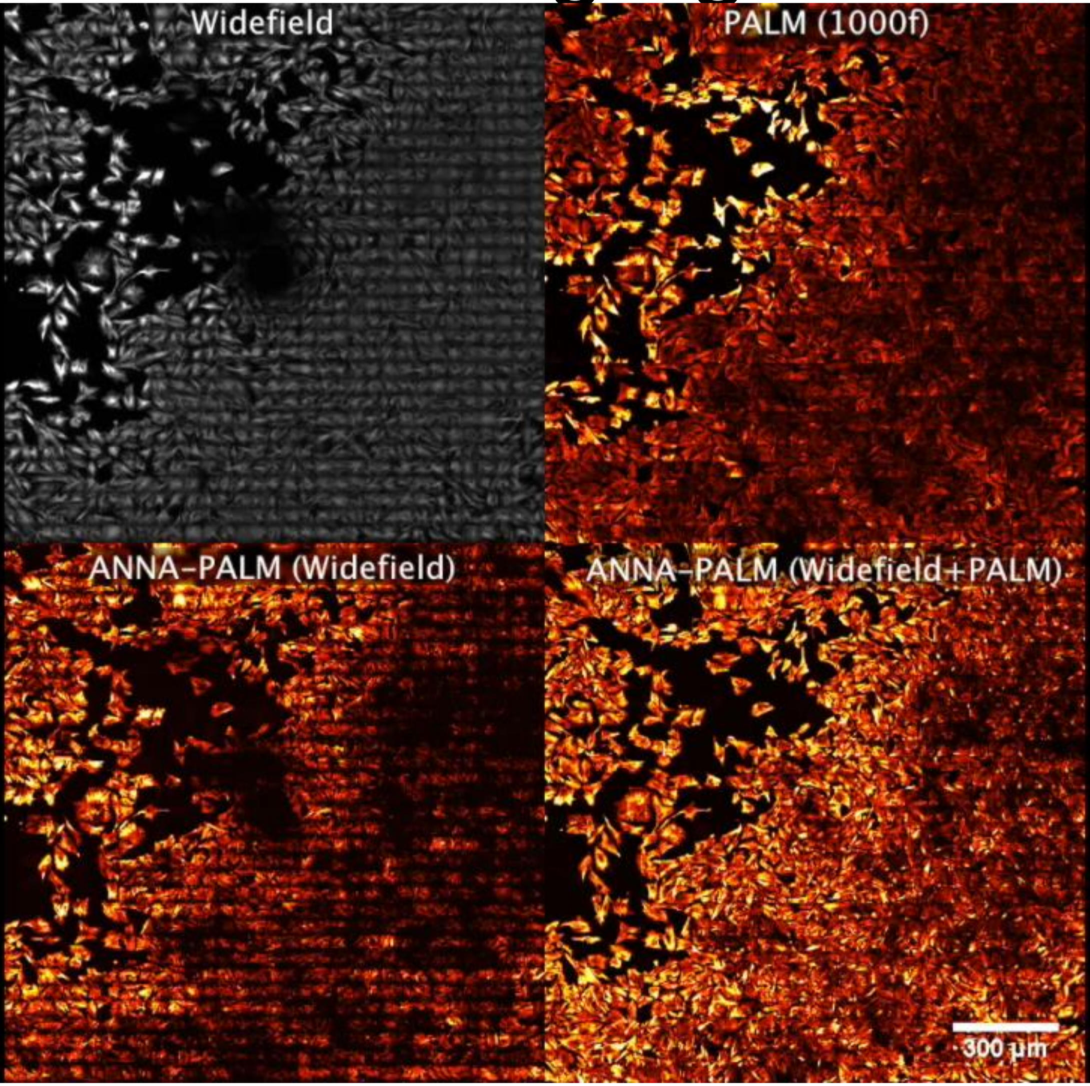

Ouyang et al. (2018)

Robustness to perturbations

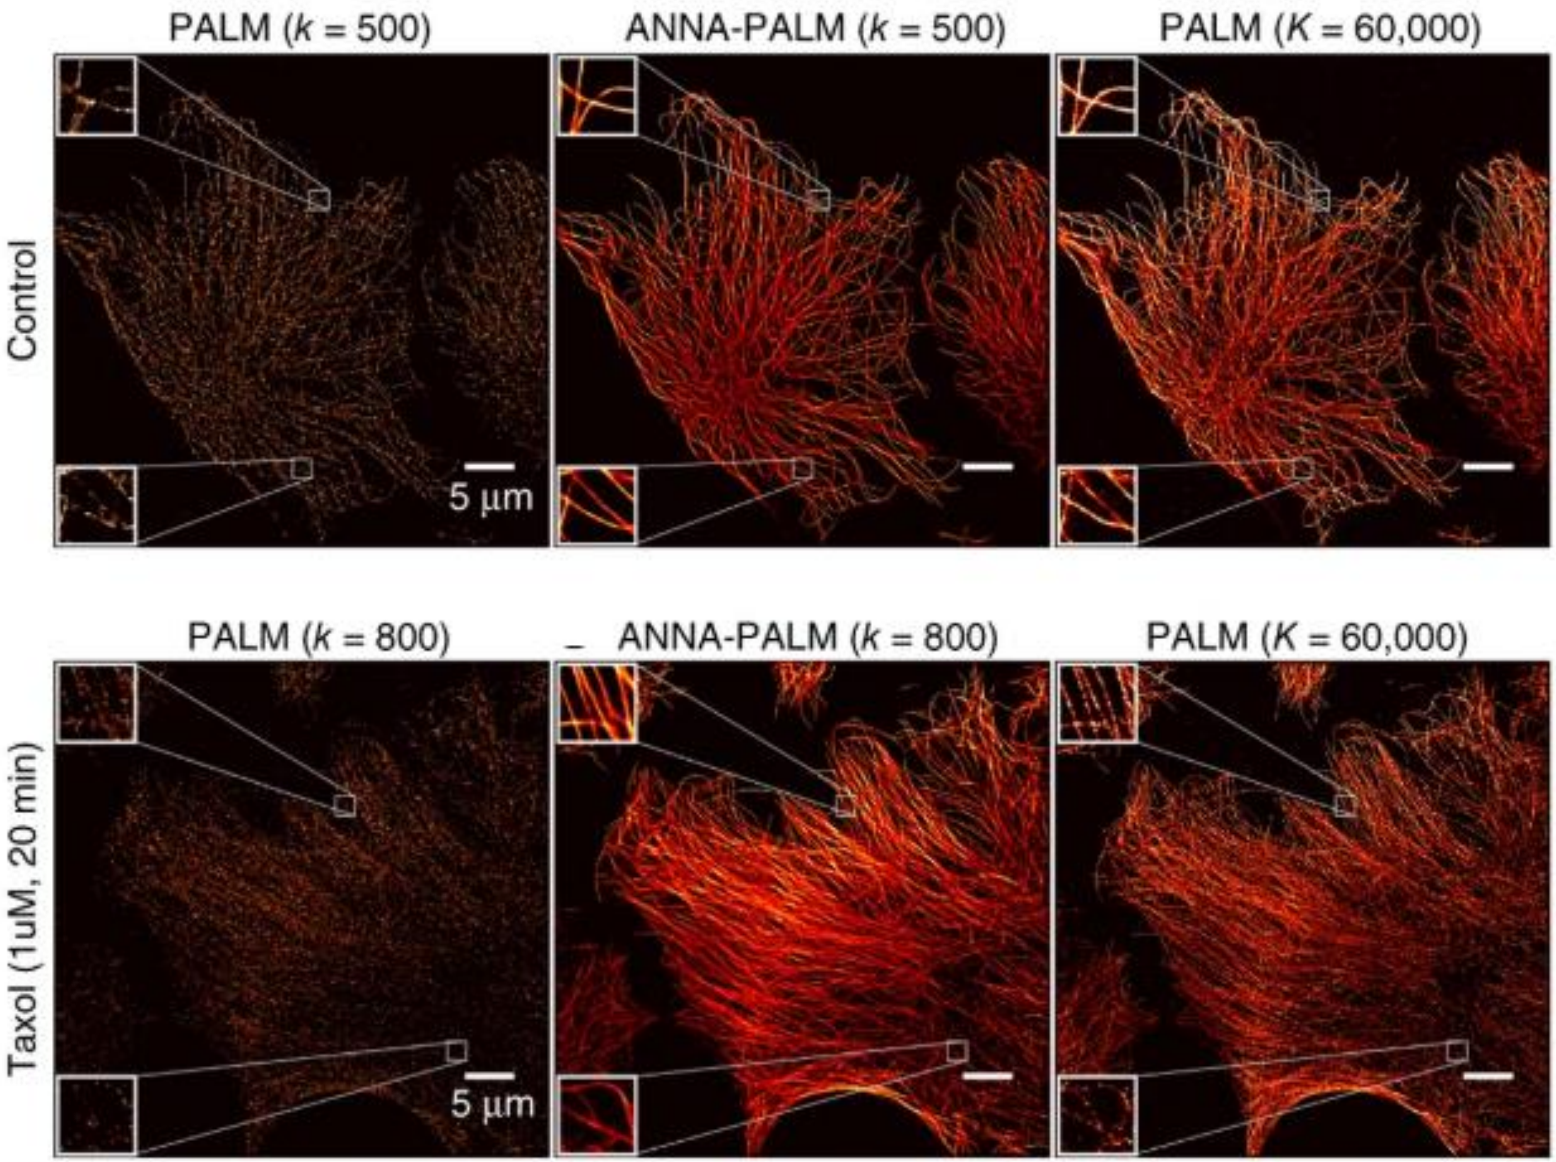

Ouyang et al. (2018)

Training with microtubules **and** (one) nuclear pores (imaged here)

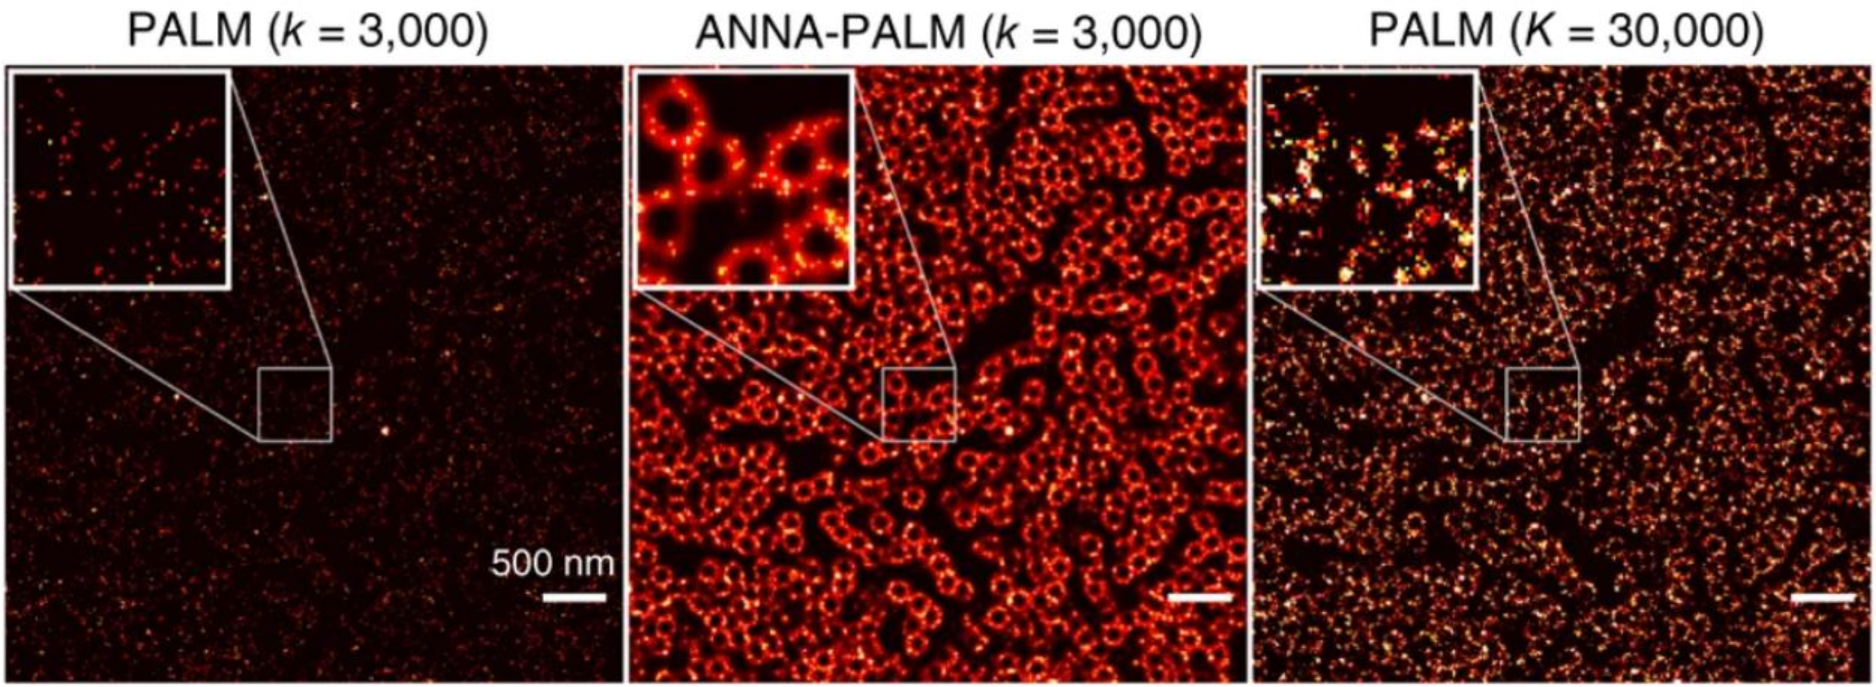

Ouyang et al. (2018)

Deep STORM
(the Israeli connection!)

Research Article Vol. 5, No. 4 / April 2018 / *Optica* 458

optica

Deep-STORM: super-resolution single-molecule microscopy by deep learning

ELIAS NEHME,[1,2] LUCIEN E. WEISS,[2] TOMER MICHAELI,[1] AND YOAV SHECHTMAN[2,*]

[1]*Electrical Engineering Department, Technion, 32000 Haifa, Israel*
[2]*Biomedical Engineering Department, Technion, 32000 Haifa, Israel*
[*]*Corresponding author: yoavsh@bm.technion.ac.il*

Received 13 February 2018; accepted 13 March 2018 (Doc. ID 323156); published 12 April 2018

Pick a paper/s and schedule yourself for class presentation

- Preferably from the list I distributed
- Include key + new (ones that I did not read yet) papers
- One (or more – from a common “topic”) paper
- Pair > single
- These paper are LONG! (compared to CS papers)
- Focus on the important stuff, in the context of our course: idea & methodology
- Reading tip: you can (mostly) ignore many biological/experimental details.
    - Example: specific molecule names. Very important, but less in the context of our course.

Course projects