ScaDS.AI
DRESDEN LEIPZIG

CENTER FOR SCALABLE DATA ANALYTICS AND ARTIFICIAL INTELLIGENCE

Deep Learning for Bio-image Analysis

Robert Haase

Funded by

Bundesministerium für Bildung und Forschung

SACHSEN

Diese Maßnahme wird gefördert durch die Bundesregierung aufgrund eines Beschlusses des Deutschen Bundestages. Diese Maßnahme wird mitfinanziert durch Steuermittel auf der Grundlage des von den Abgeordneten des Sächsischen Landtags beschlossenen Haushaltes.

Machine learning for image analysis

In classical machine learning, we typically select features for training our classifier

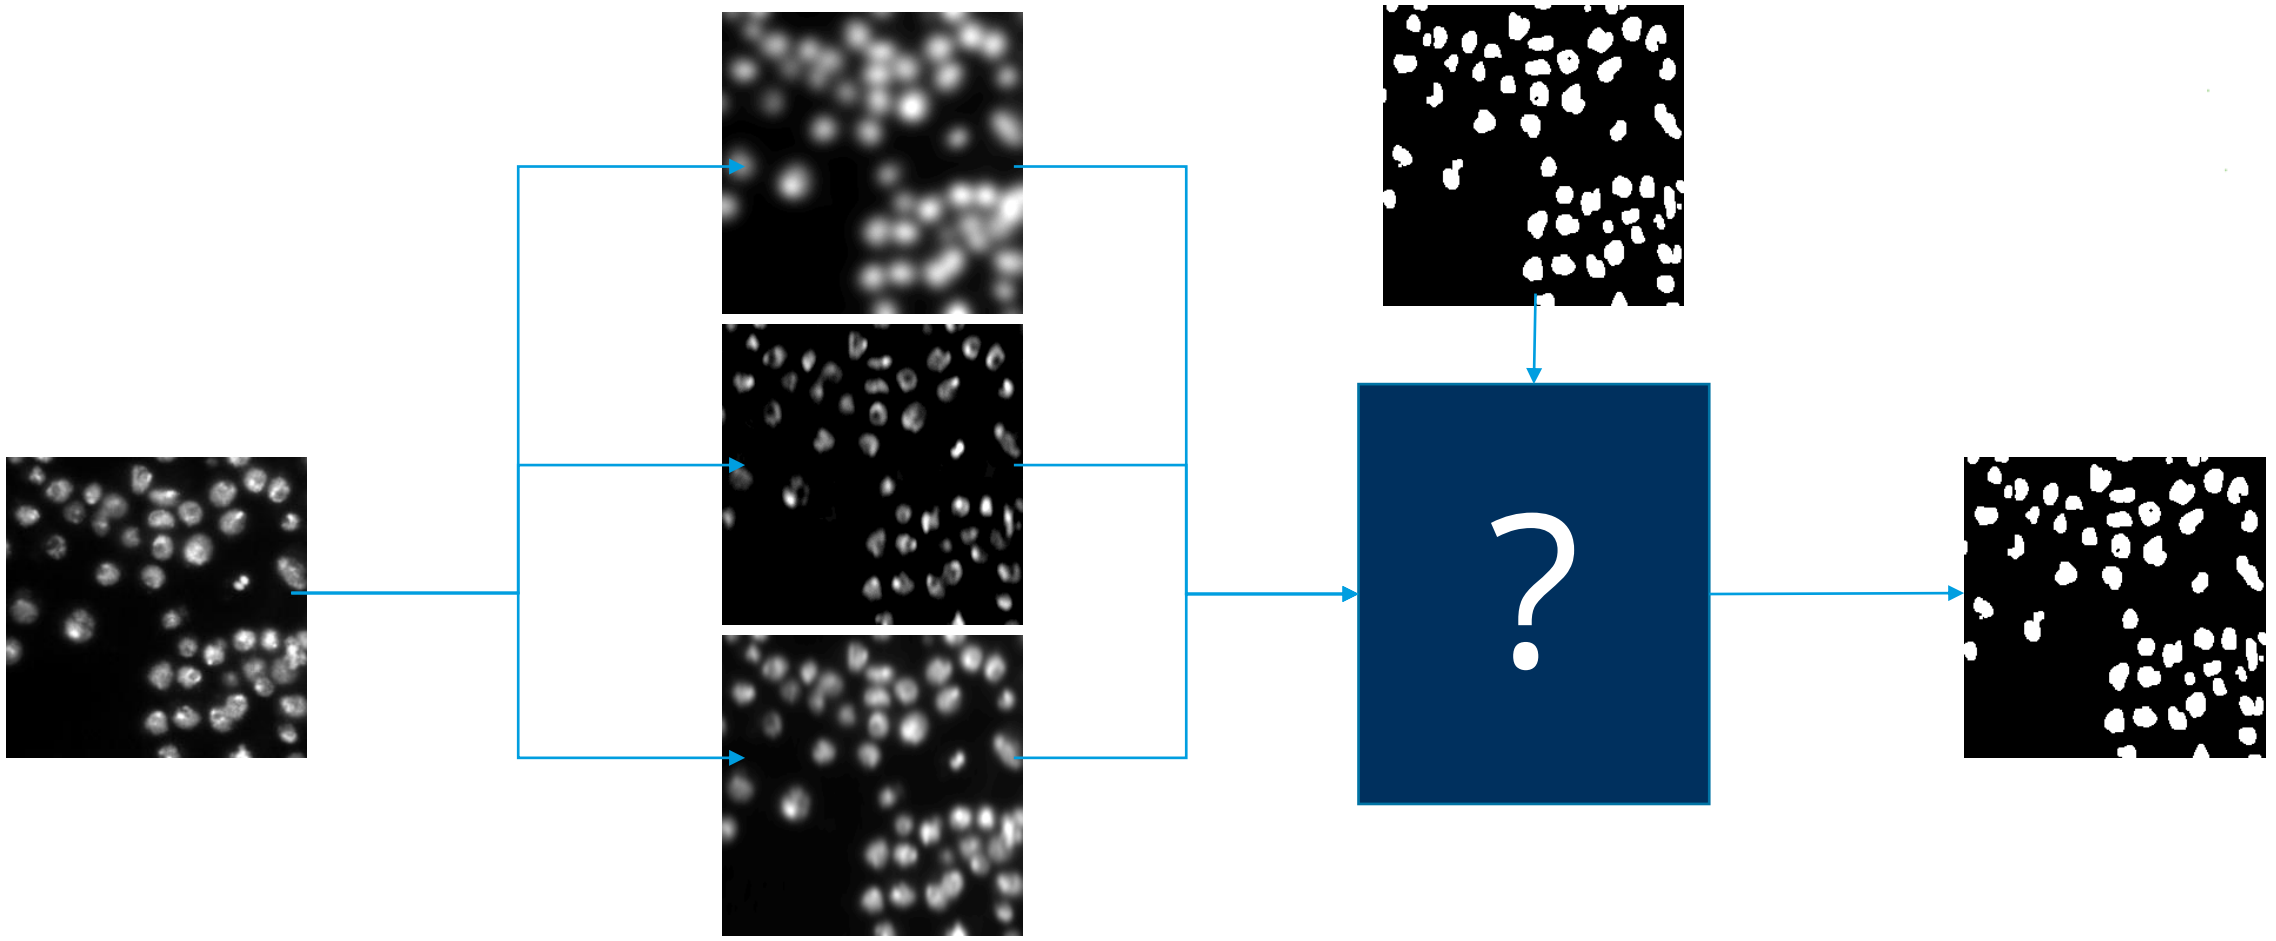

Image data source: BBBC038v1, available from the Broad Bioimage Benchmark Collection (Caicedo et al., Nature Methods, 2019)

Outlook: Deep learning for image analysis

In deep learning, this selection becomes part of the black box

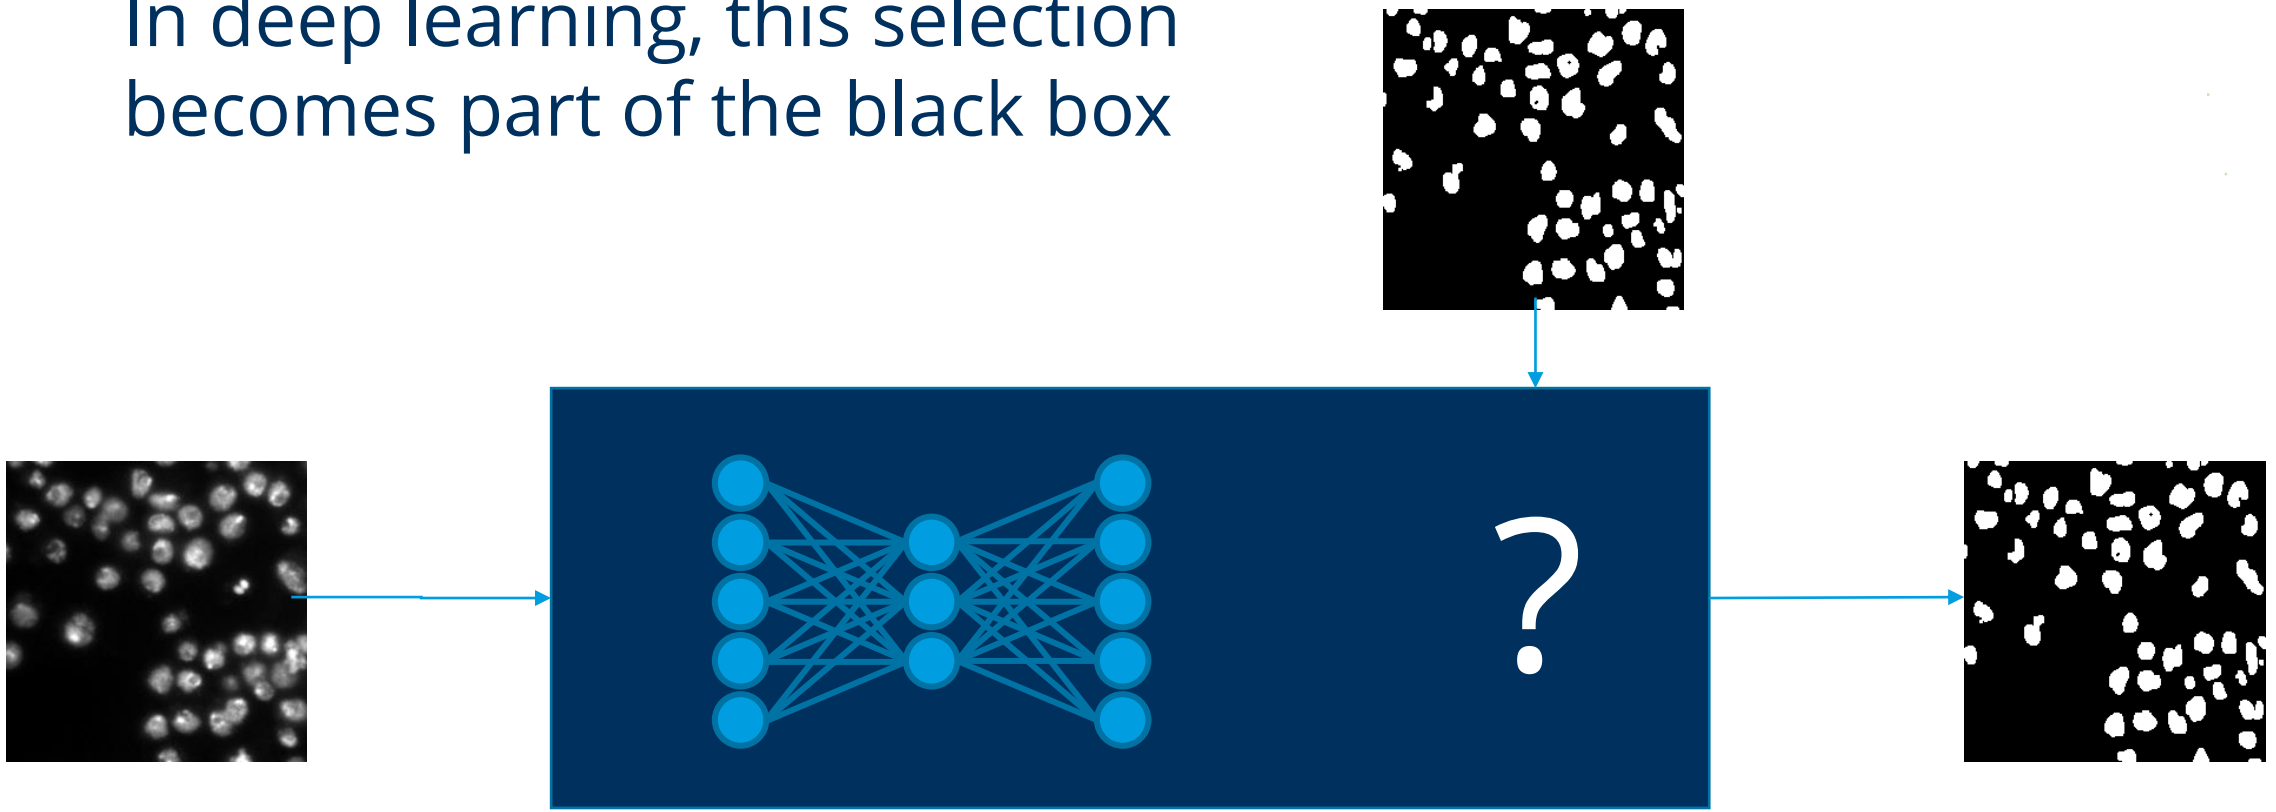

Convolutional neural networks

Image data source: BBBC038v1, available from the Broad Bioimage Benchmark Collection (Caicedo et al., Nature Methods, 2019)

Neural networks

- How biologists see neurons

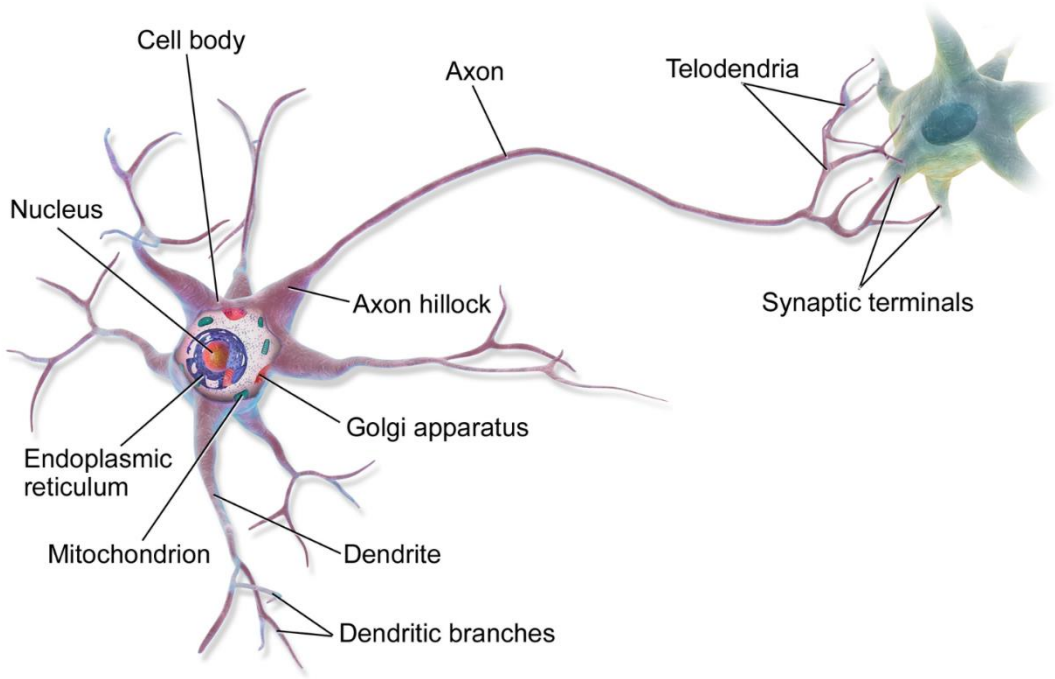

- How computer scientists see neurons

"perceptron"

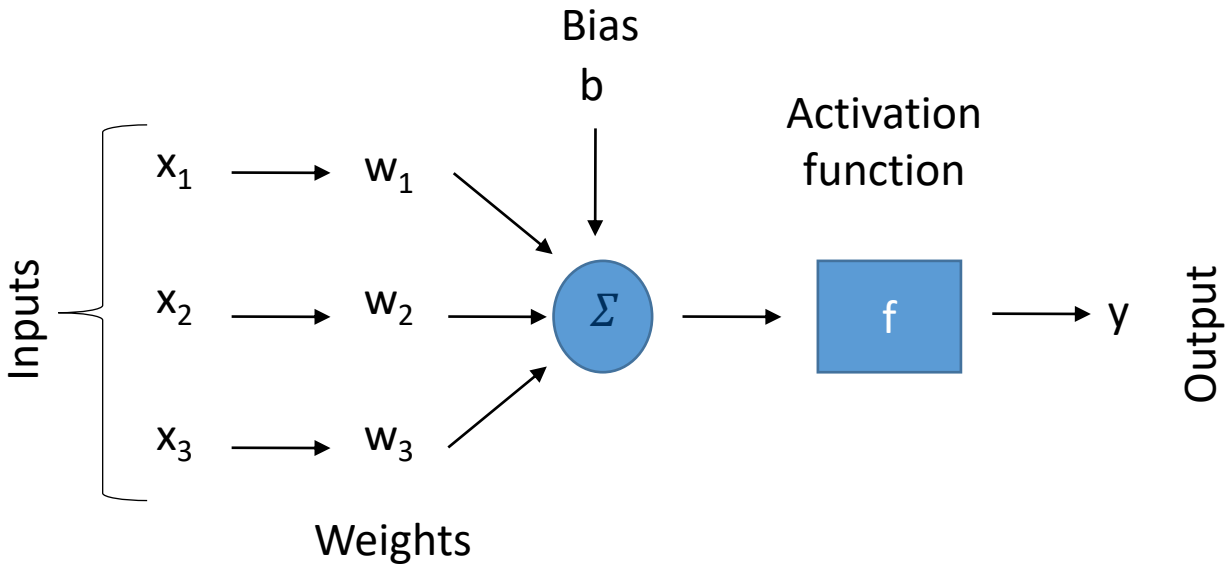

$$y = f(w|\ \ |1x_1 + w_2x_2 + w_3x_3 + b)$$

Neuron image source:
https://commons.wikimedia.org/wiki/File:Blausen_0657_MultipolarNeuron.png

Neural Networks

- Early form: “Multilayer Perceptron”
- fully connected class of feedforward artificial neural network

If there are *many* hidden layers, we speak of a *deep* neural network

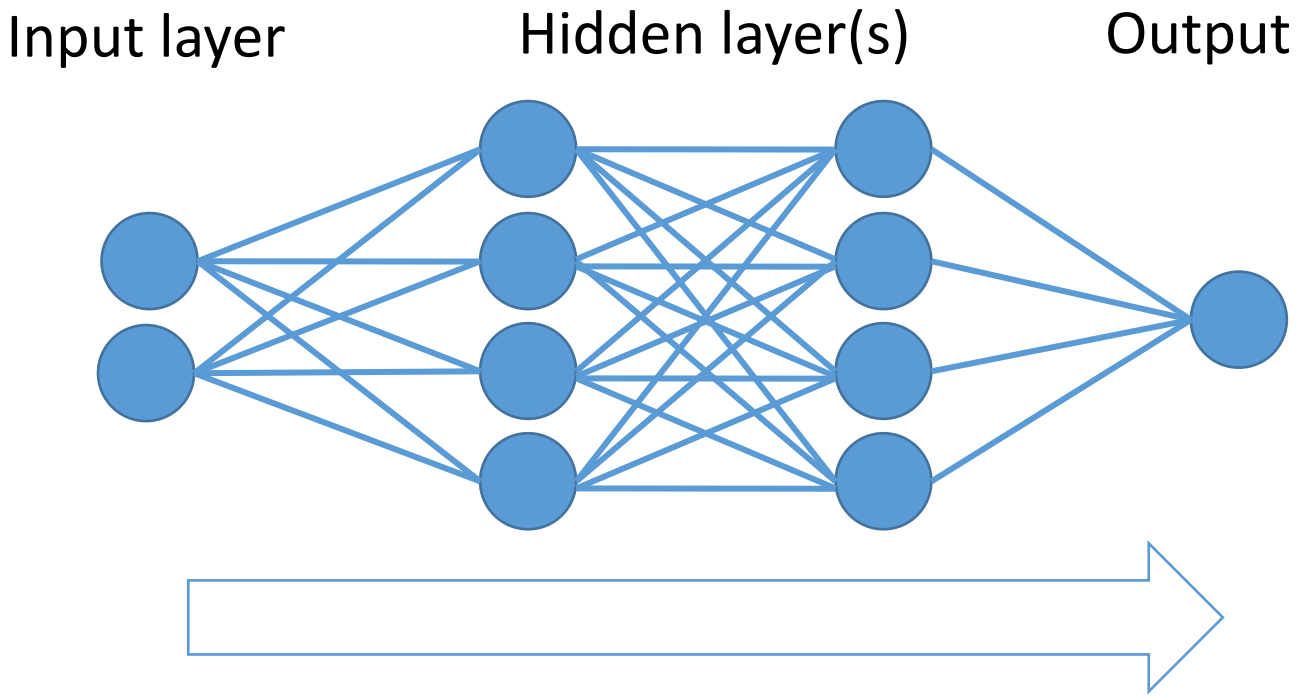

Convolutional neural networks

- Layer types

Fully connected layer

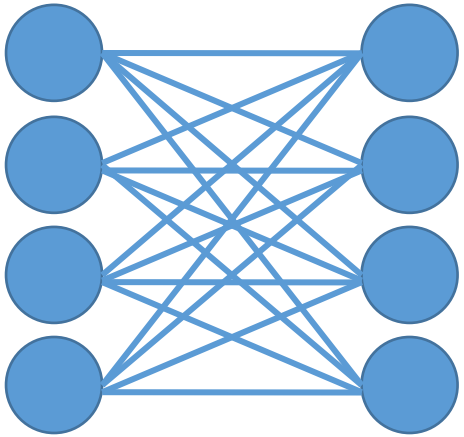

Convolutional neural networks

- Layer types

Fully connected layer

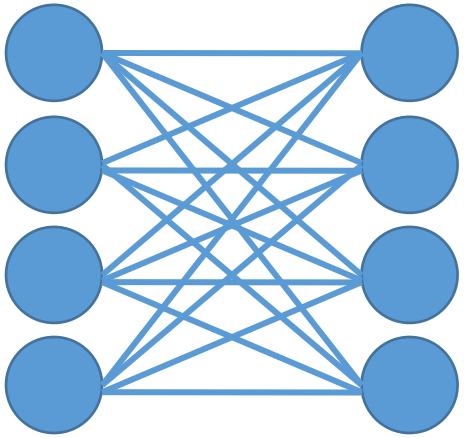

Convolutional layer

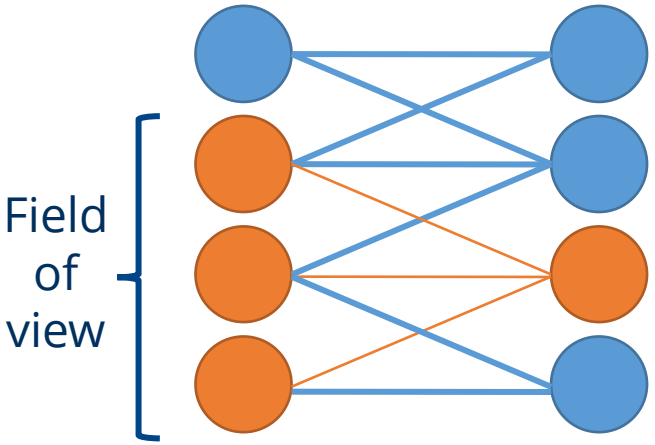

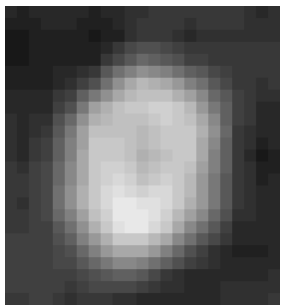

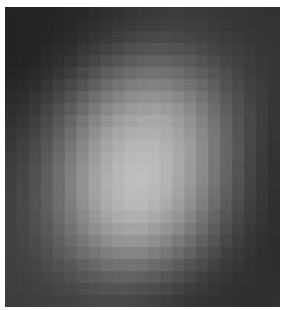

Deep Learning for Microscopy

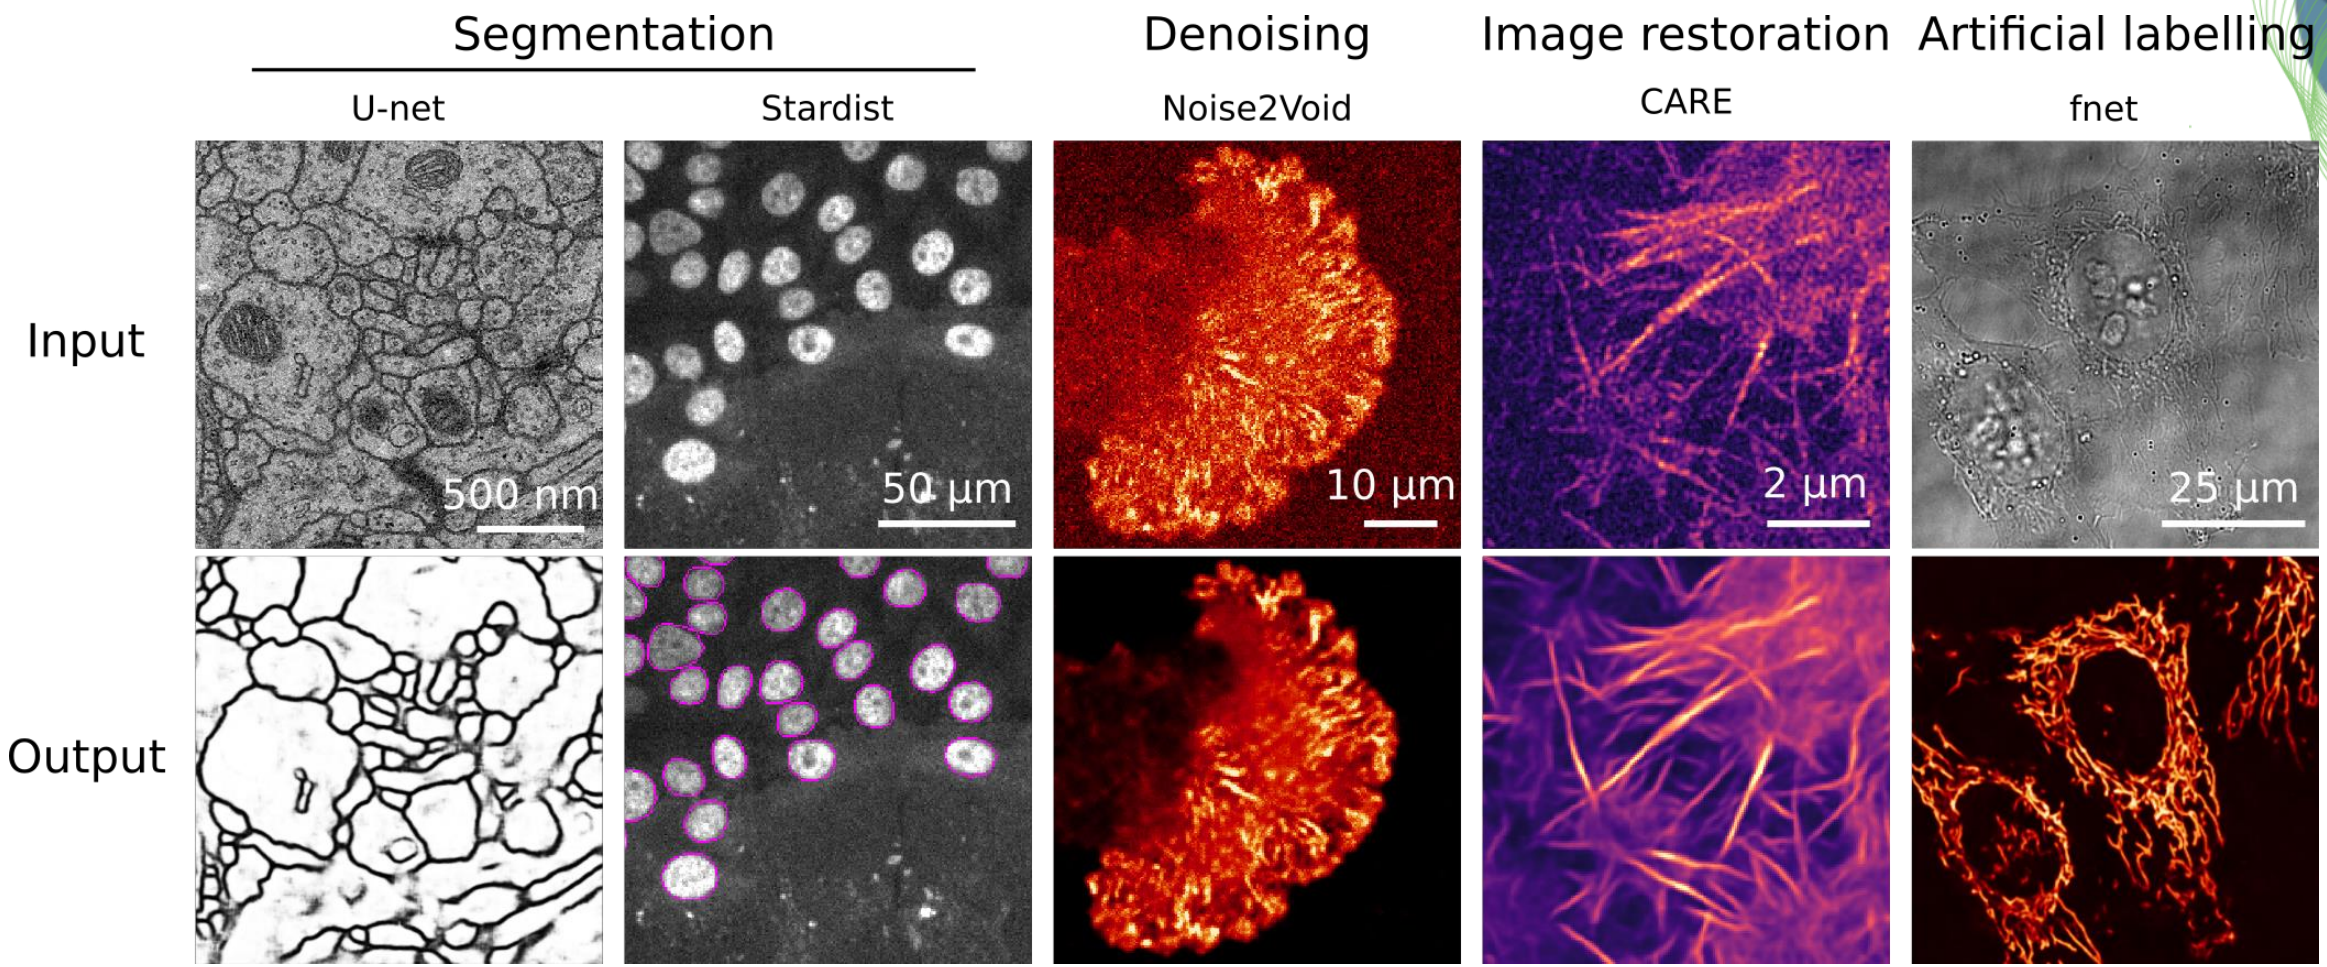

Source: Ricardo Henriques lab, https://github.com/HenriquesLab/ZeroCostDL4Mic/blob/master/Wiki_files/FigureS2.png license: MIT

Trade-offs in live-cell imaging

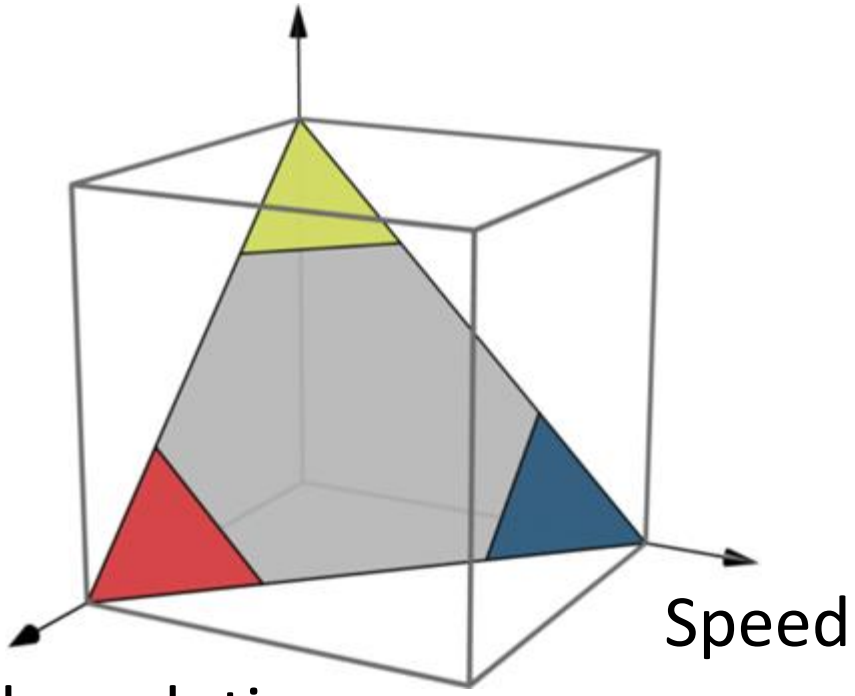

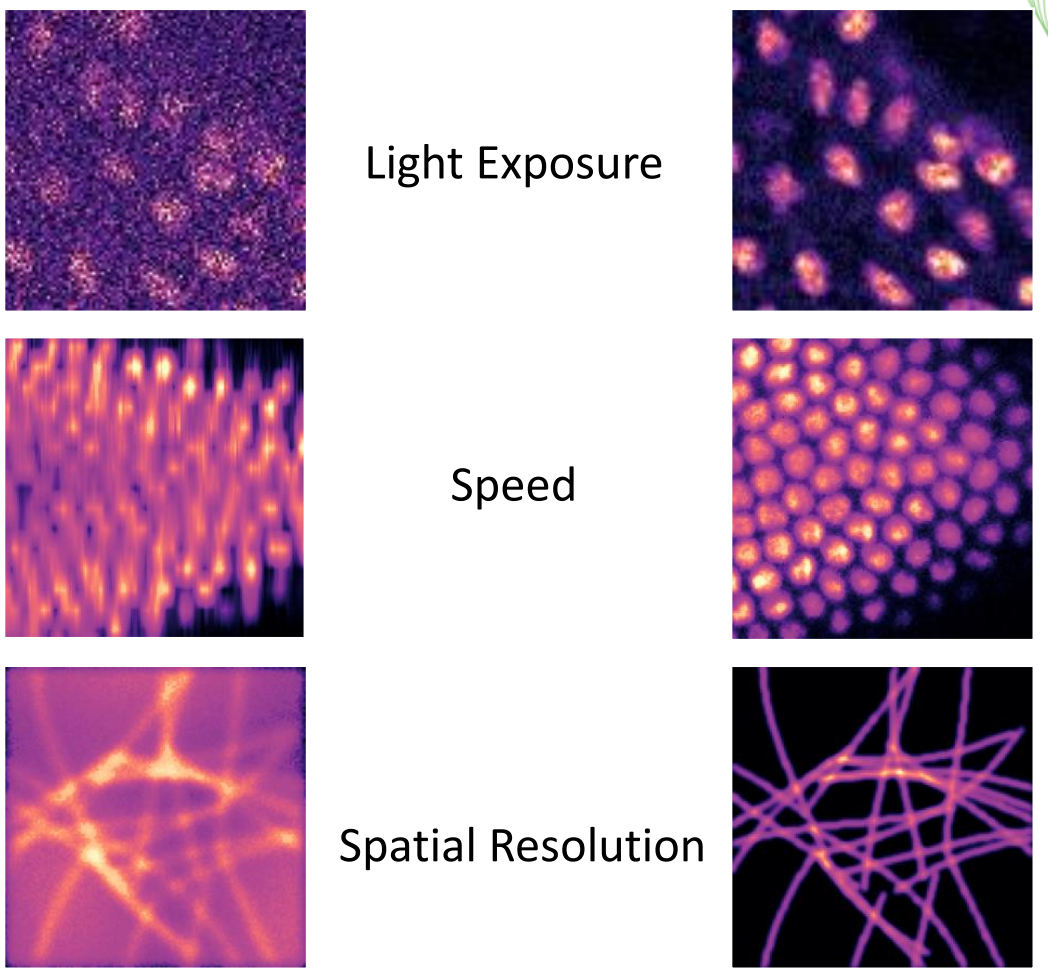

The U-net

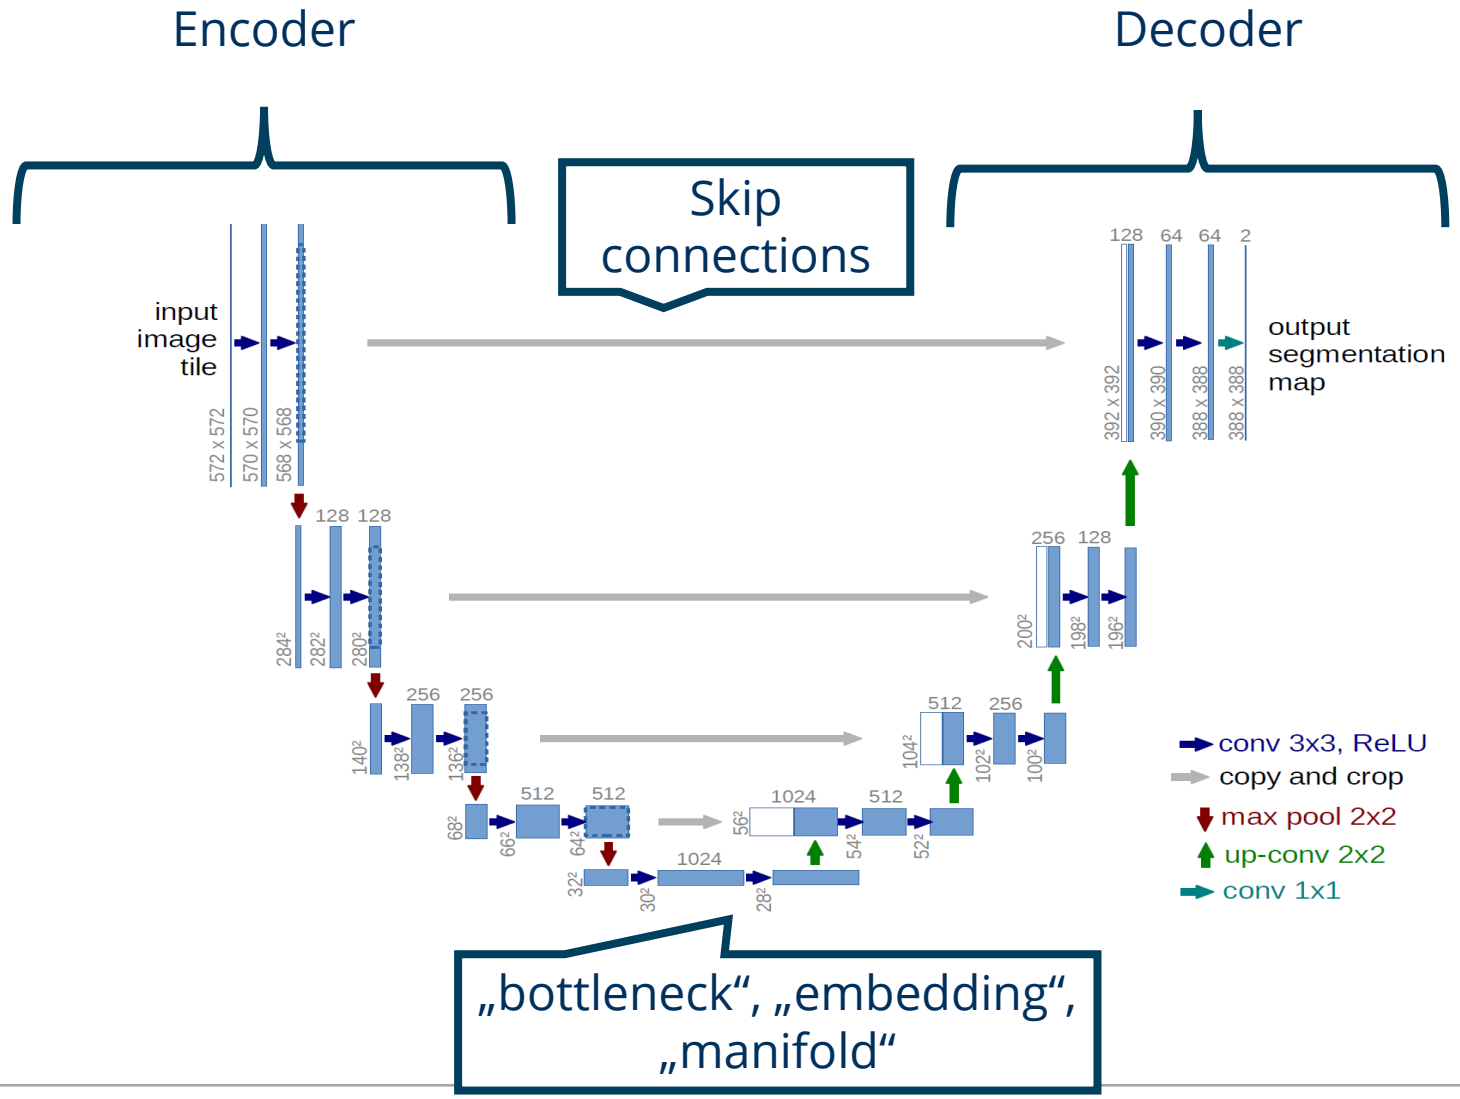

Source: Ronneberger (2015)
https://arxiv.org/pdf/1505.04597.pdf

Image denoising: CARE

- Content aware image restoration (CARE)
- Image acquisition of pairs of images: A high-quality and a low-quality image.
- Problem: shot noise, biology moves!

Pair of images required!

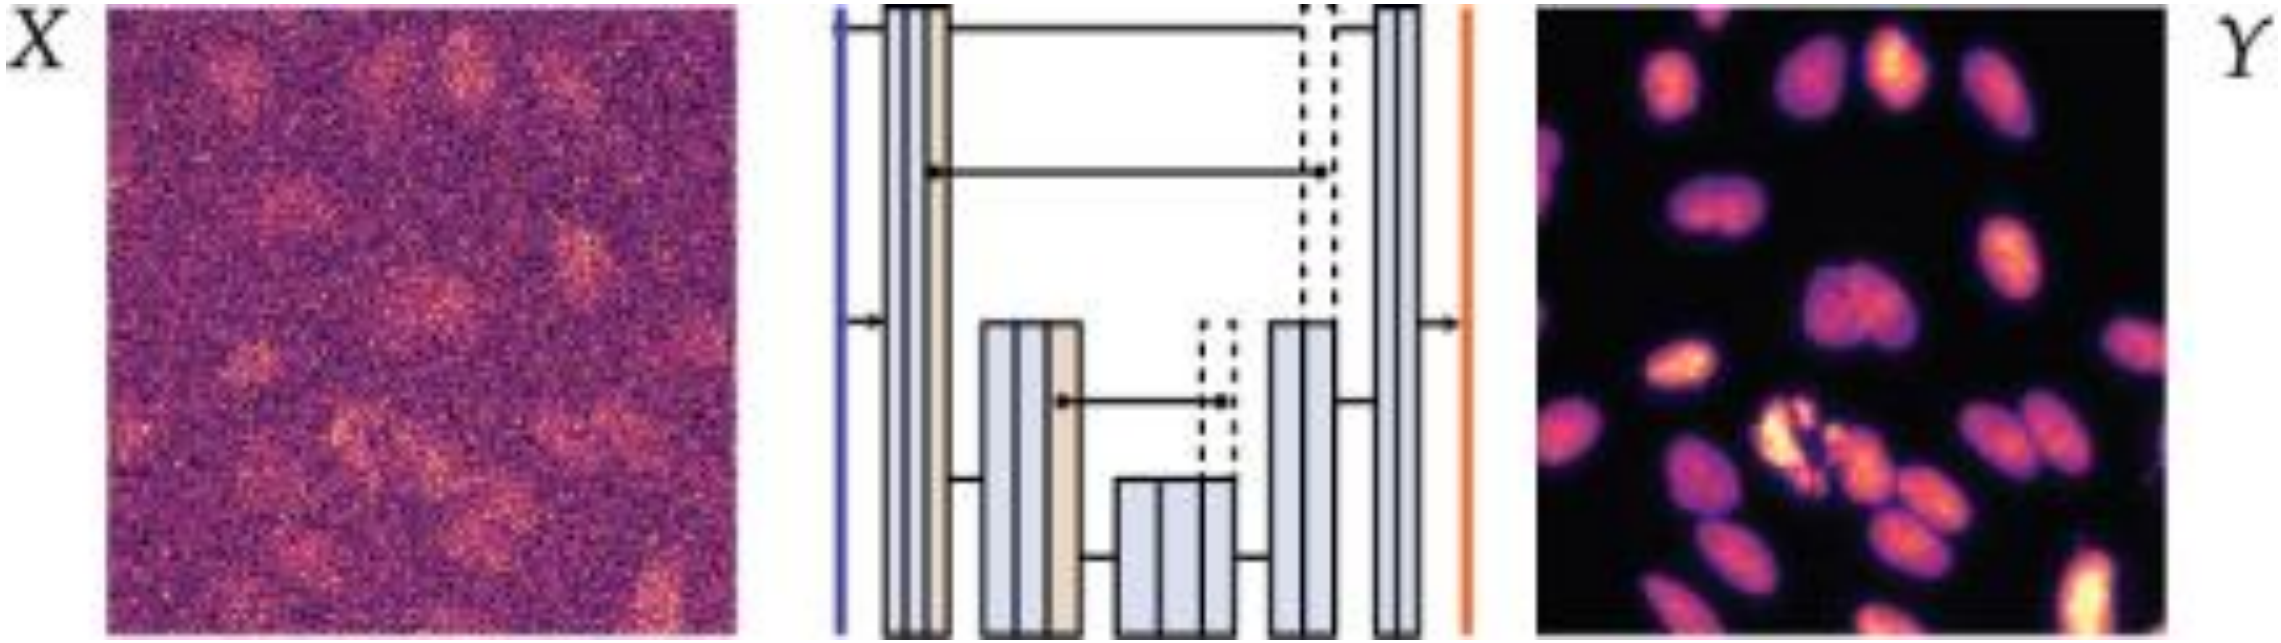

Source: https://github.com/CSBDeep/CSBDeep/blob/main/examples/denoising2D/2_training.ipynb, License: BSD3

Image denoising: CARE

- Important to use on the same conditions/structures/staining that the networks were trained on!

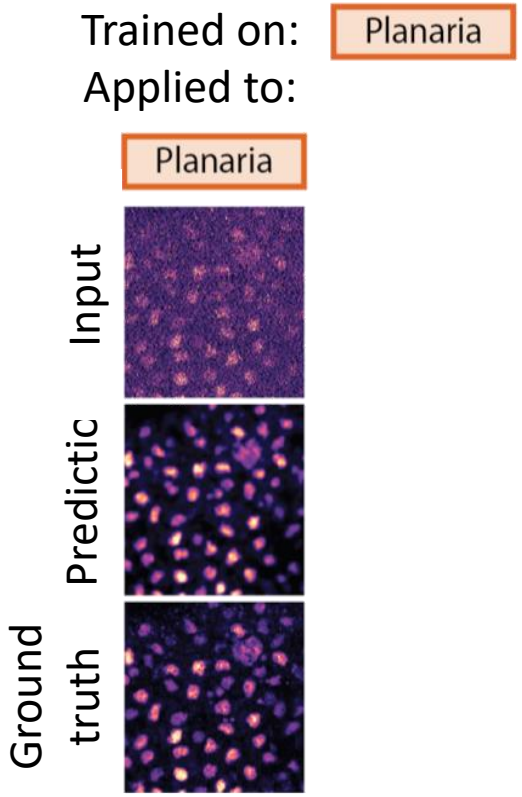

Noise2noise

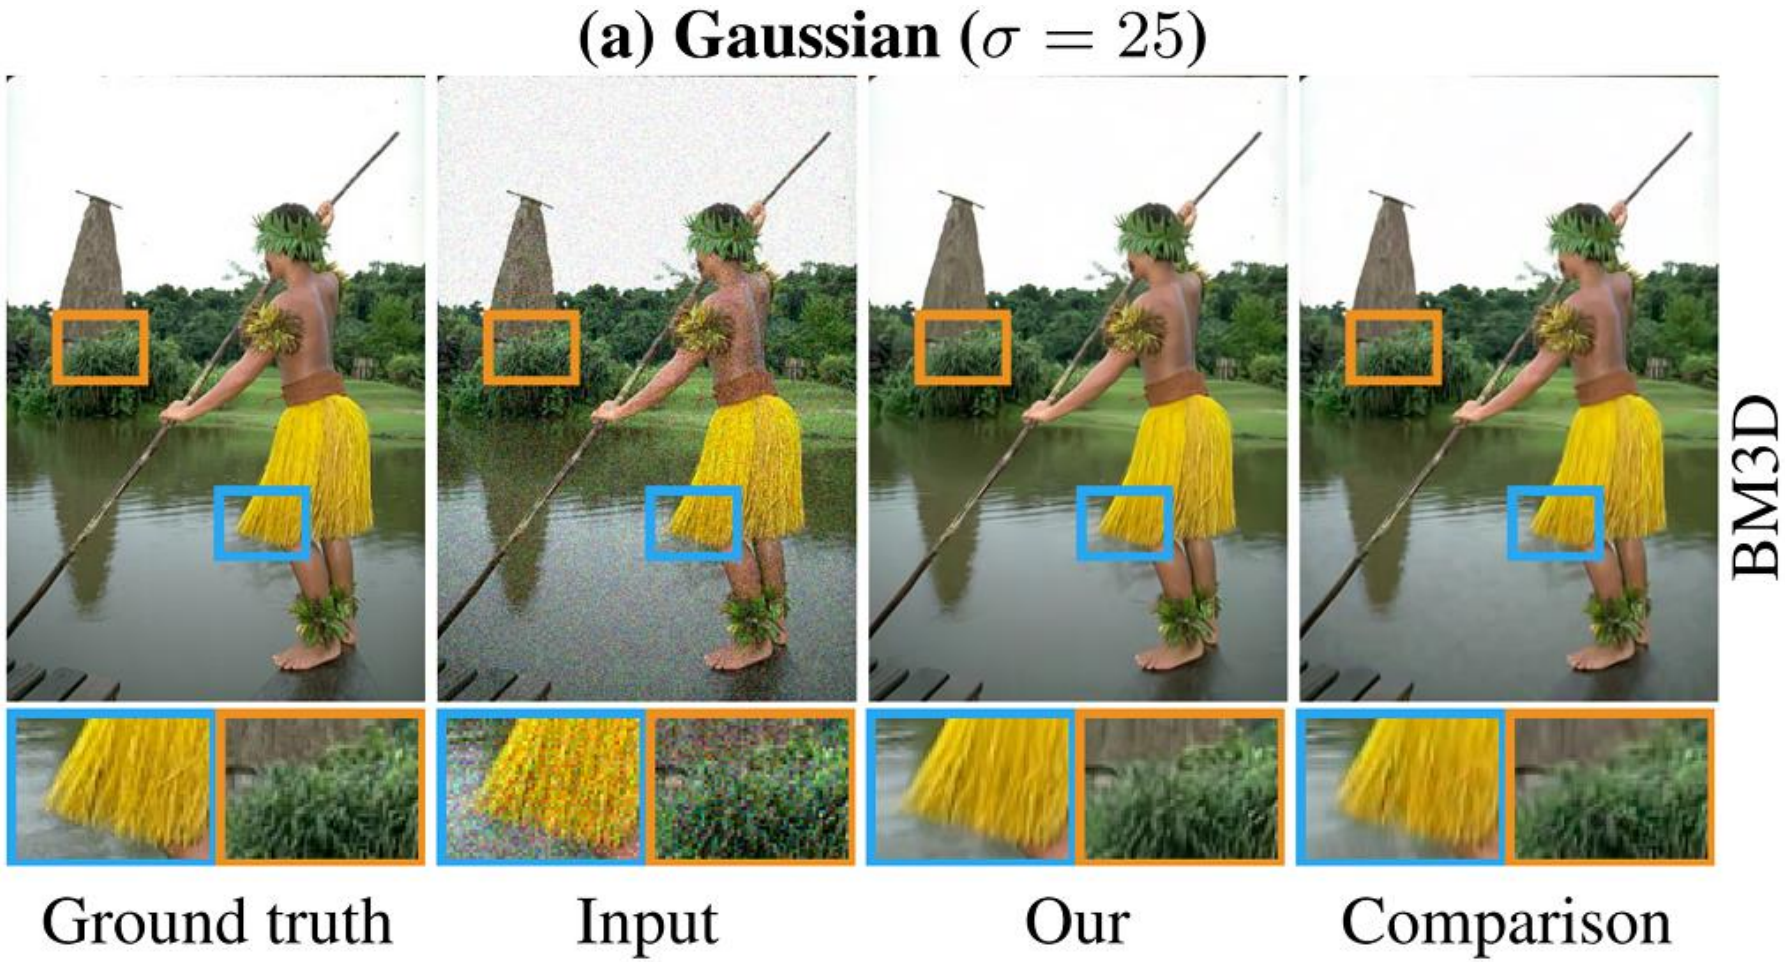

Source: Noise2Noise: Learning Image Restoration without Clean Data
Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, Timo Aila
https://arxiv.org/abs/1803.04189
https://github.com/NVlabs/noise2noise License: CC-BY-NC 4.0

Noise2noise

- by Nvidia (Lehtinen 2018) https://arxiv.org/pdf/1803.04189.pdf

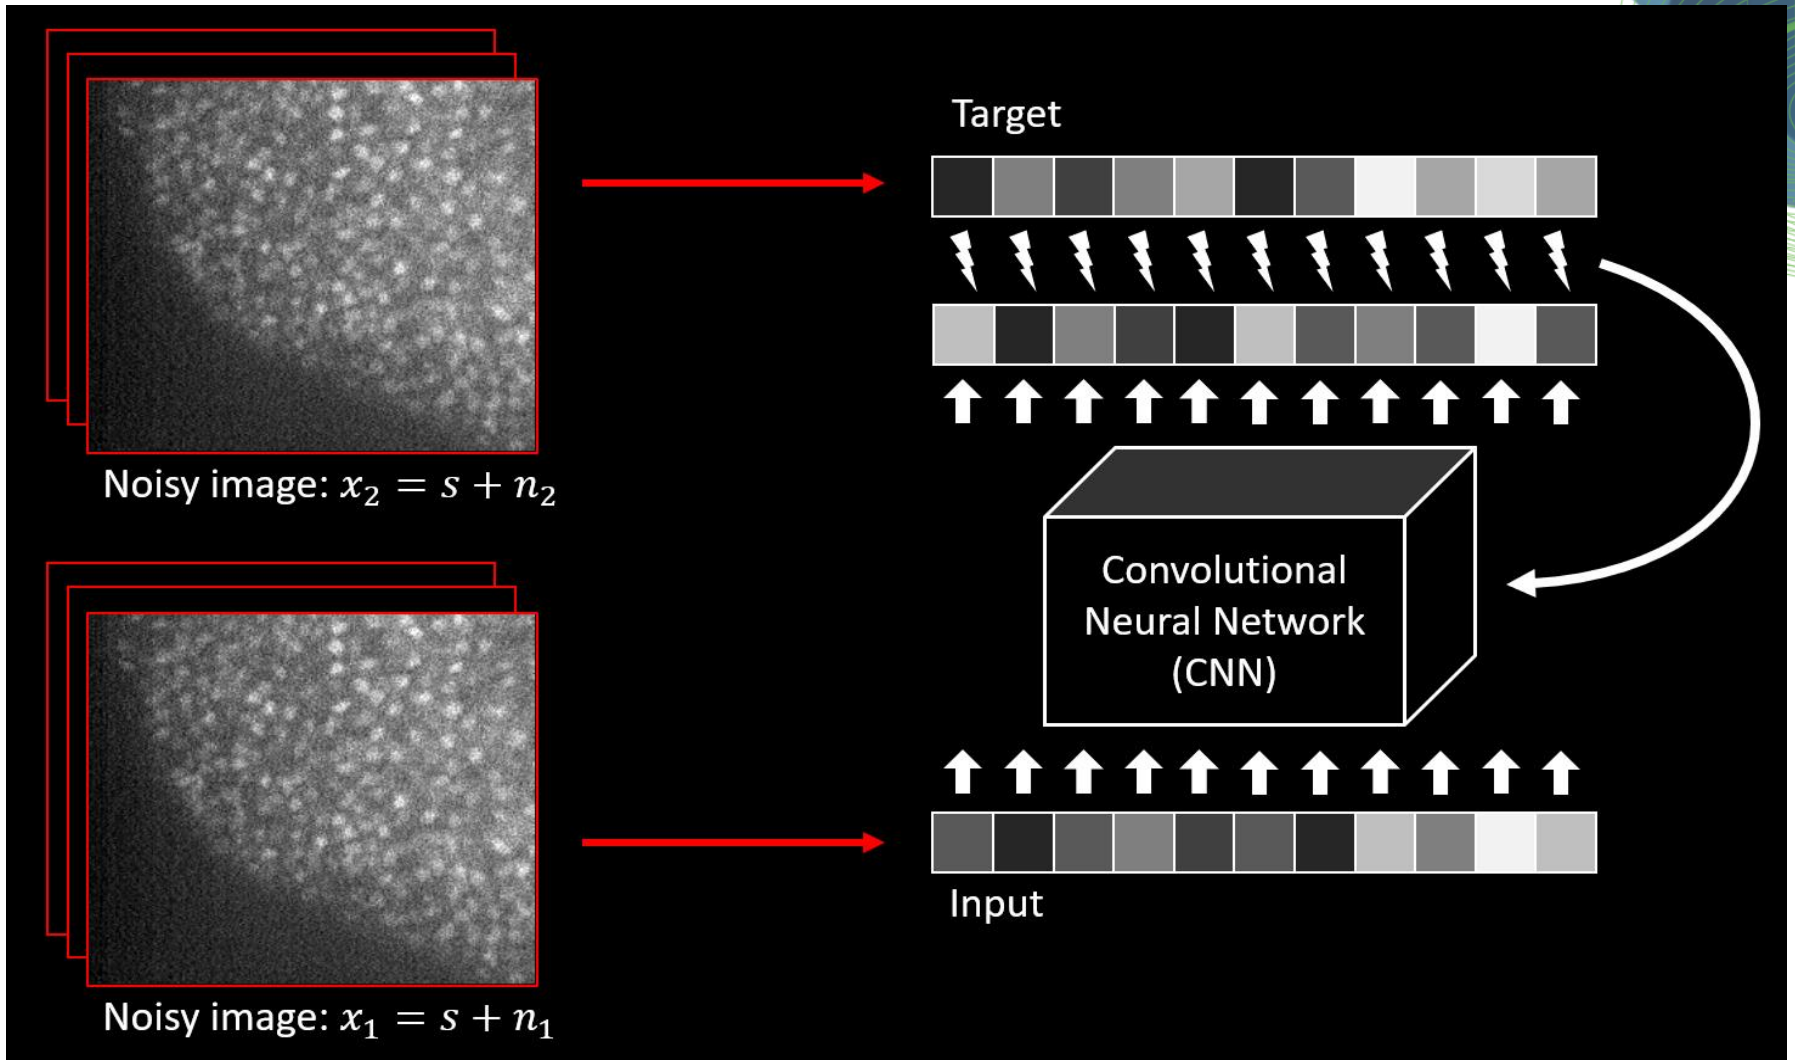

Noise2void

Image denoising without image pairs

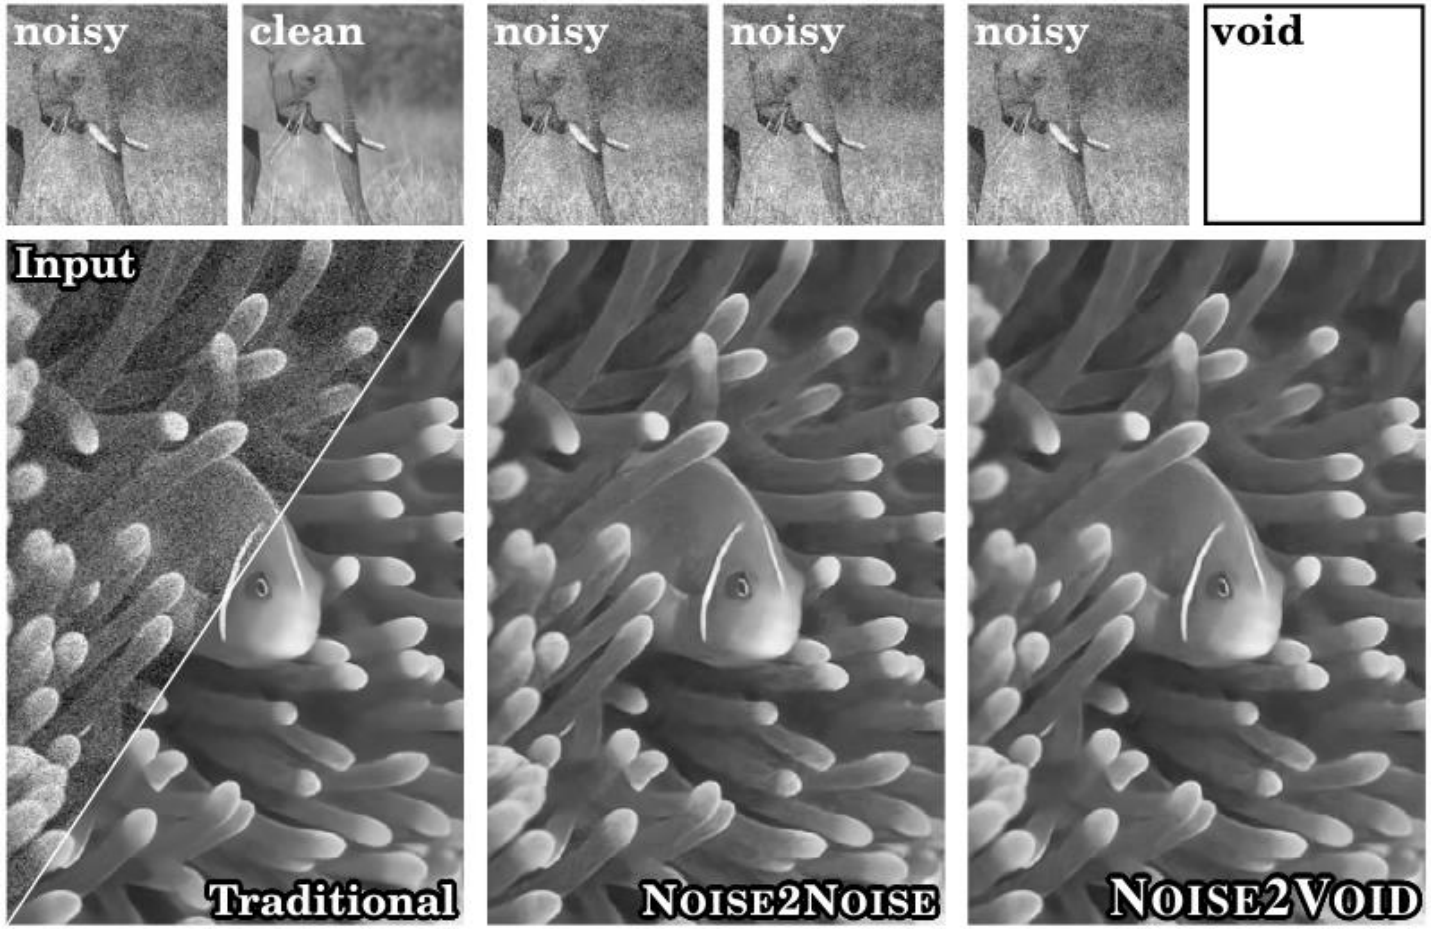

Source: A.Krull et al.
https://arxiv.org/pdf/1811.10980

Noise2void

- Noise2void: Krull et al (2019) https://arxiv.org/abs/1811.10980
- Noise2self: Batson and Royer (2019) https://arxiv.org/abs/1901.11365

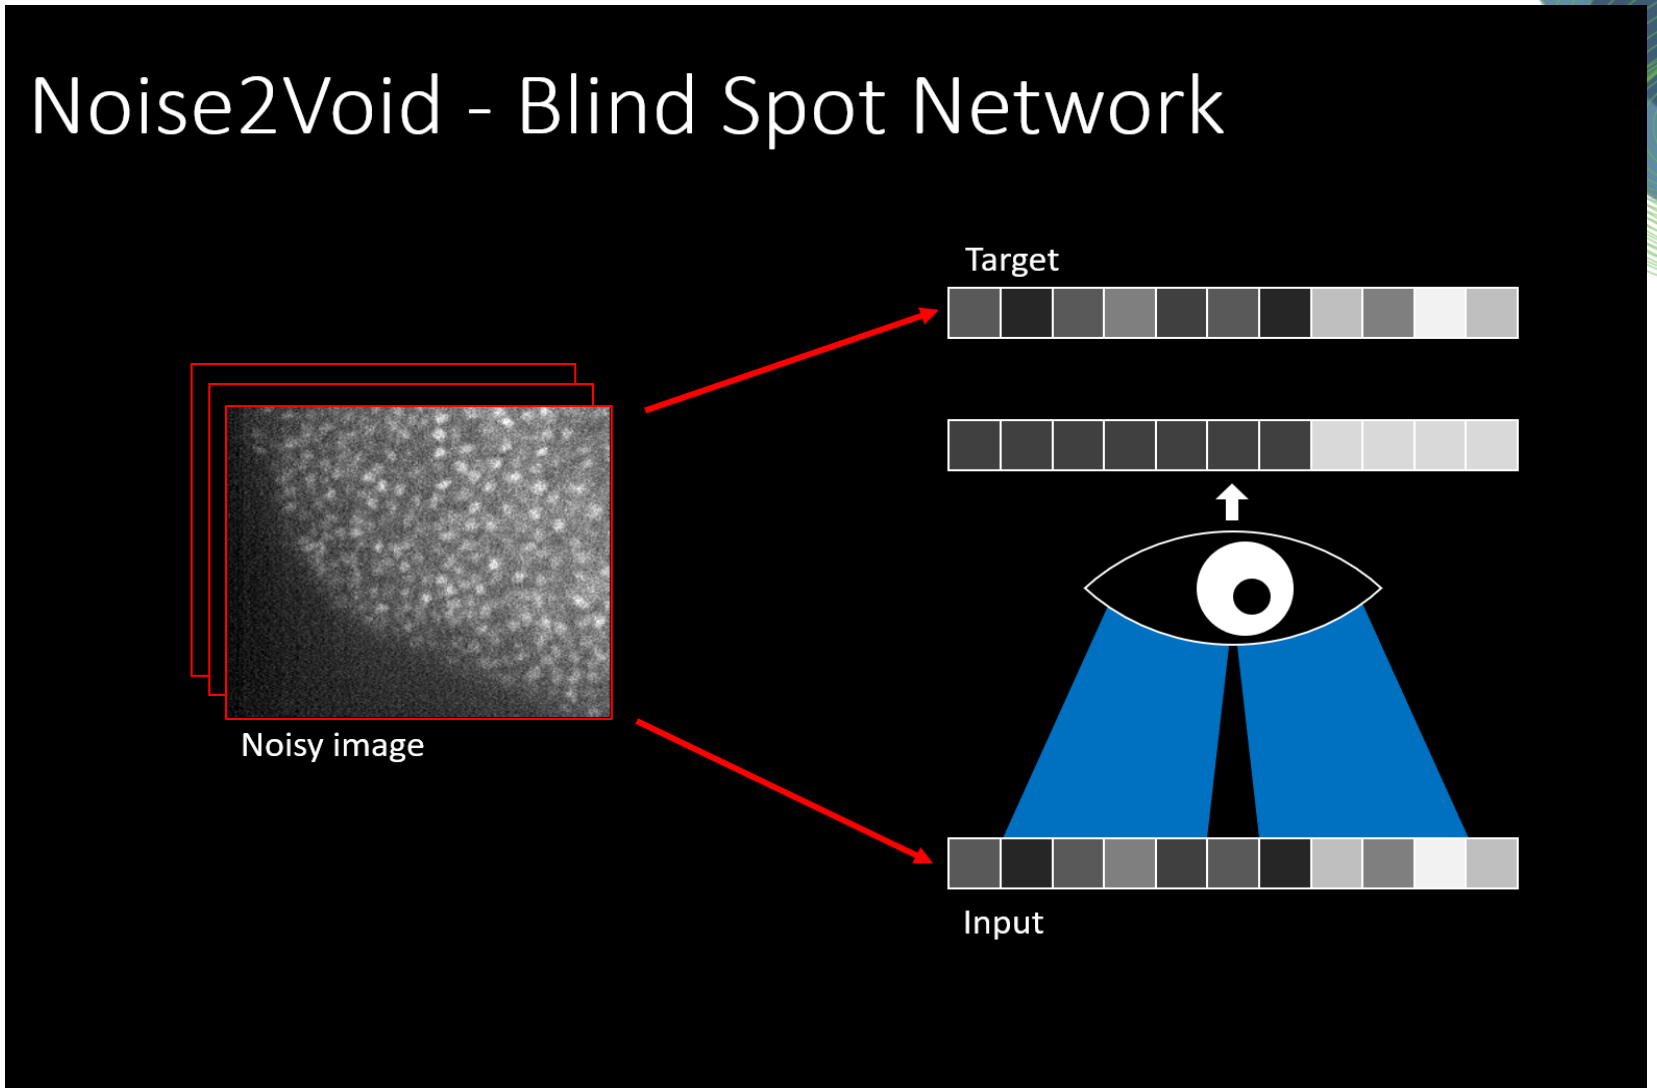

Noise2void

„Blind-spot-network“

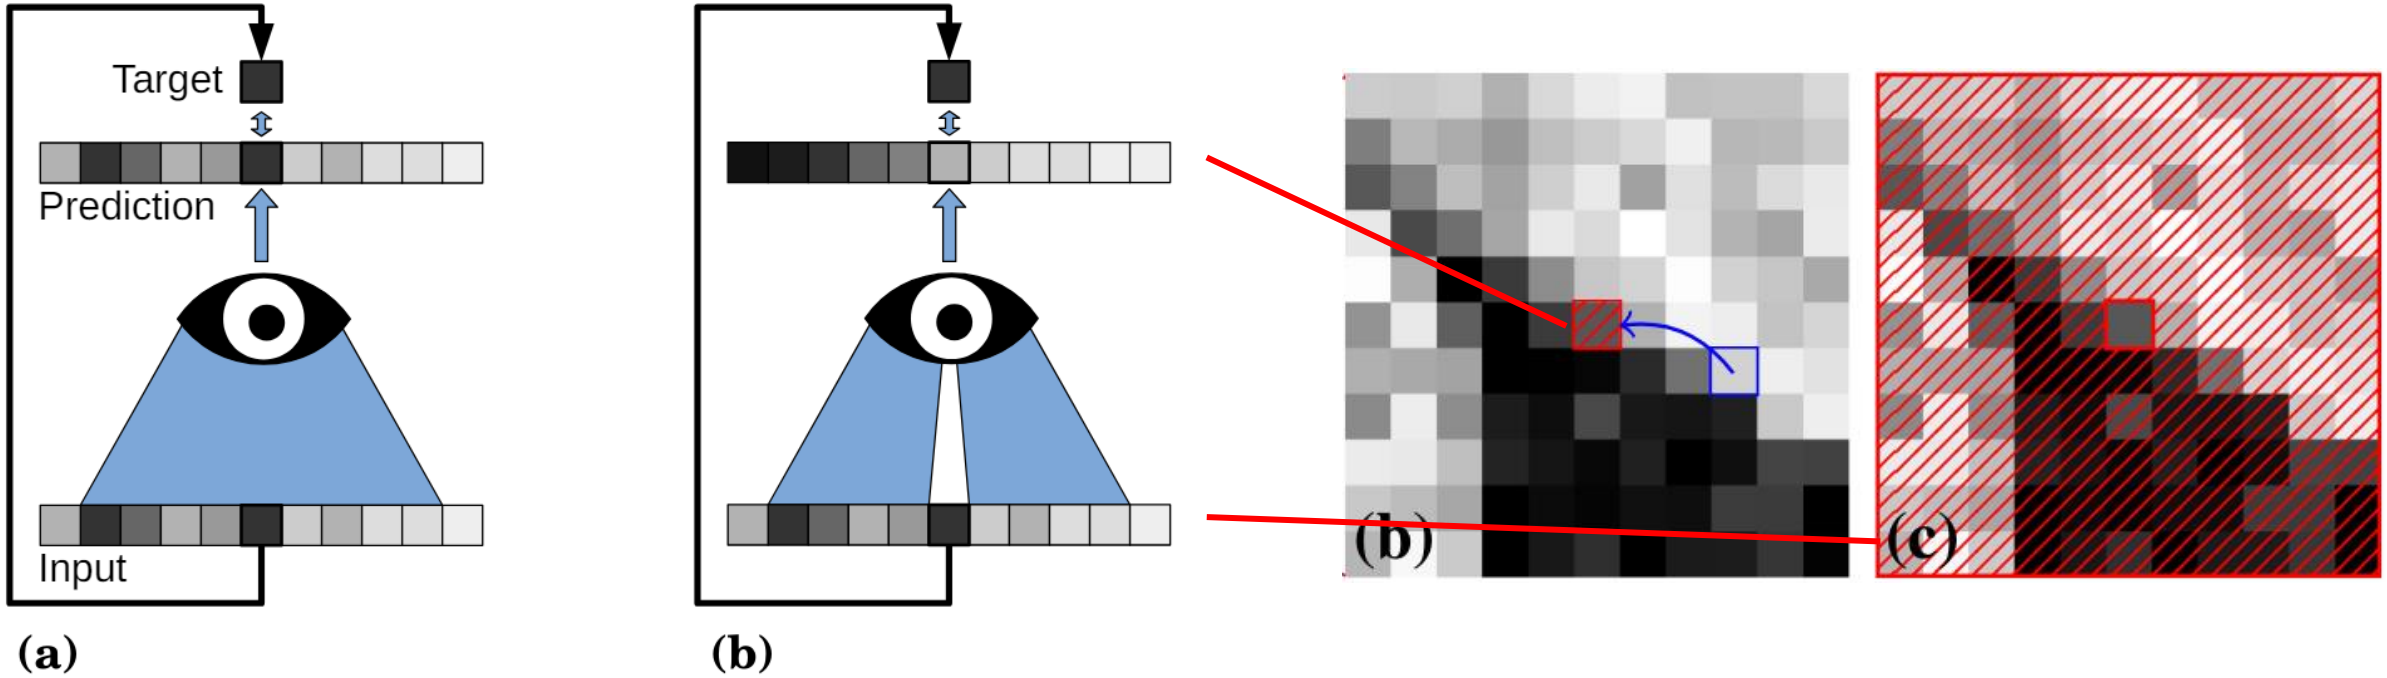

Noise2void

Napari-plugin

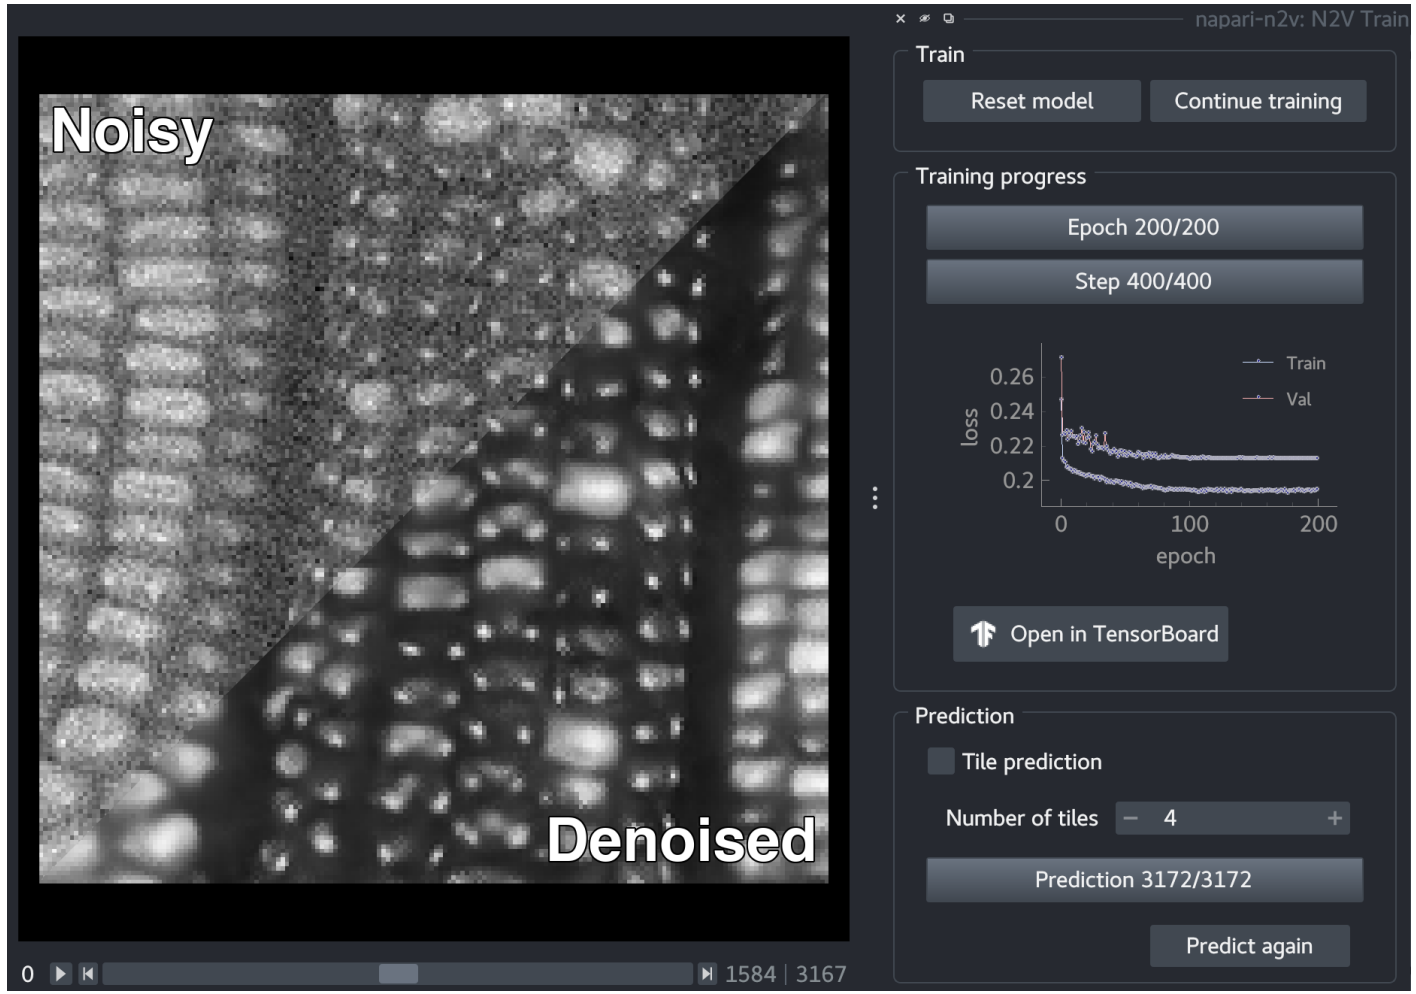

Source: https://github.com/juglab/napari-n2v

Image denoising

- Noise2Void

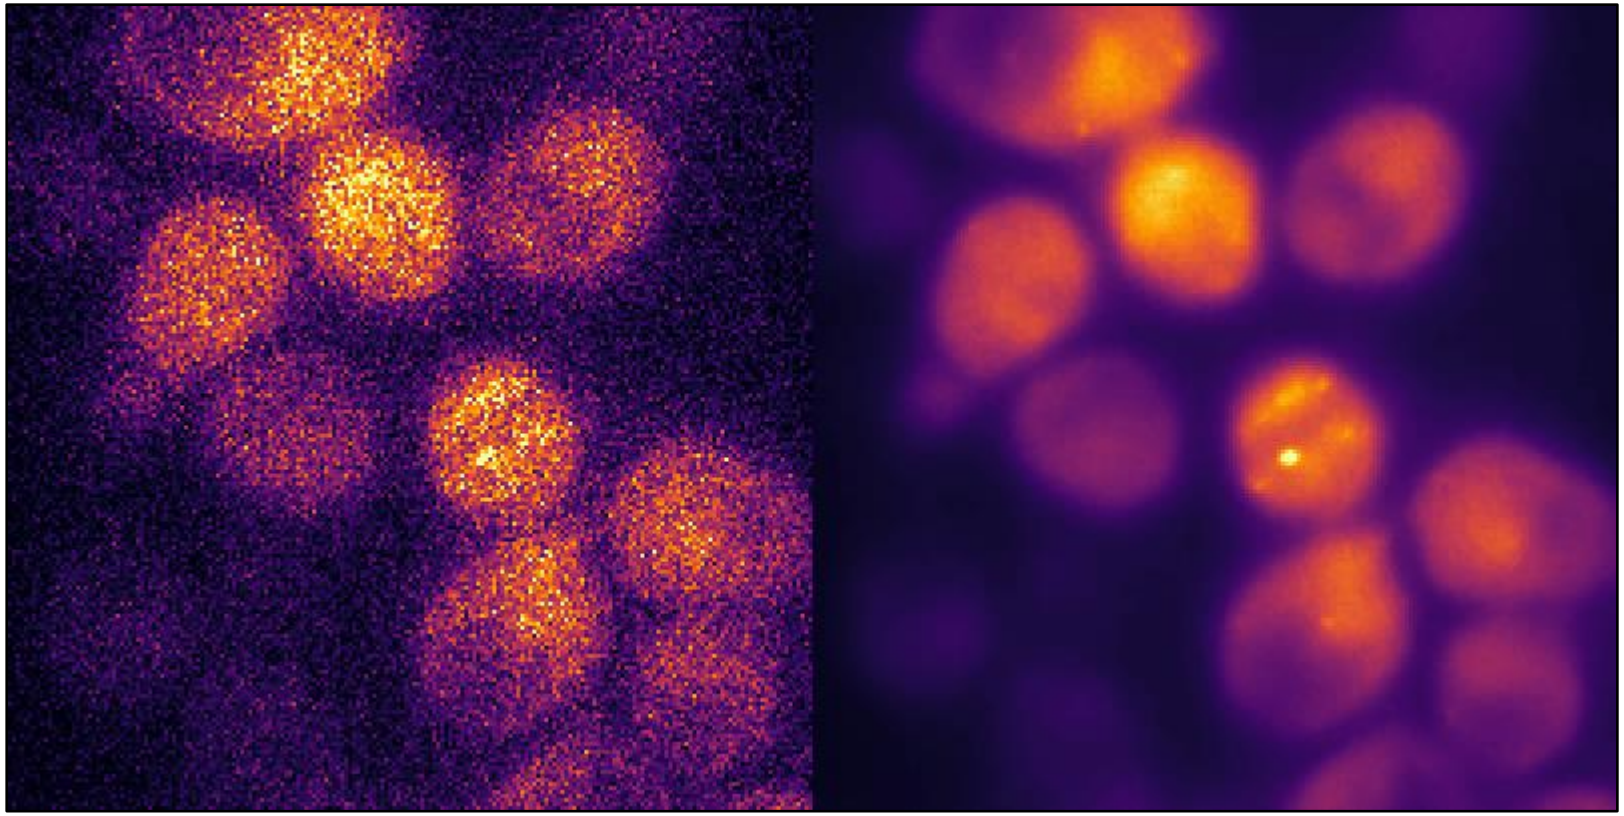

Data by Stephanie Heinrich

Slide adapted from: Alexander Krull, MPI CBG

StarDist: Nuclei segmentation

Advanced algorithms are necessary when nuclei become too dense.

Image of nuclei

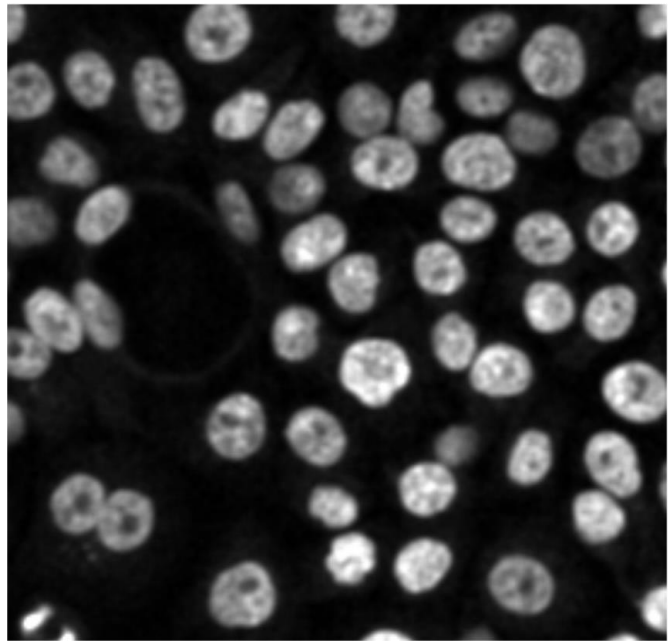

Corresponding masks

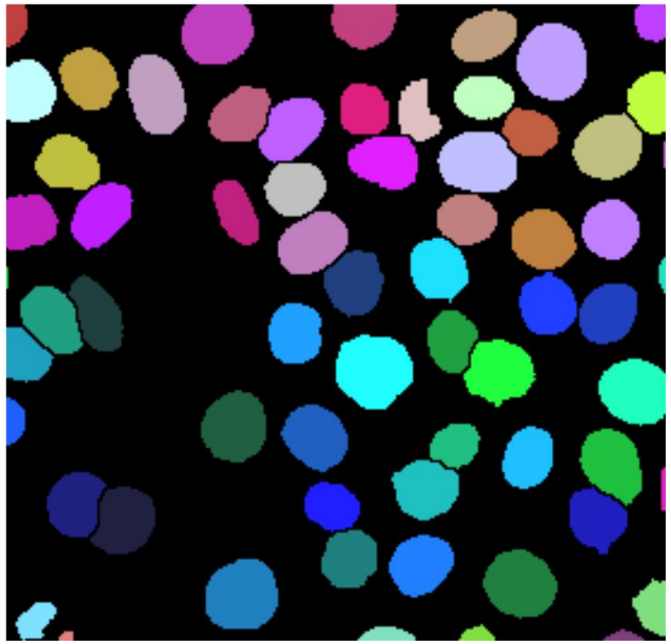

Source: Ricardo Henriques lab, https://github.com/HenriquesLab/ZeroCostDL4Mic/blob/master/Wiki_files/Stardist_nuclei_masks.png license: MIT

StarDist

- Prediction of probable object centers + polygon outlines
- Non-maximum-suppression of less likely polygons

Class. Image proc.

Deep learning

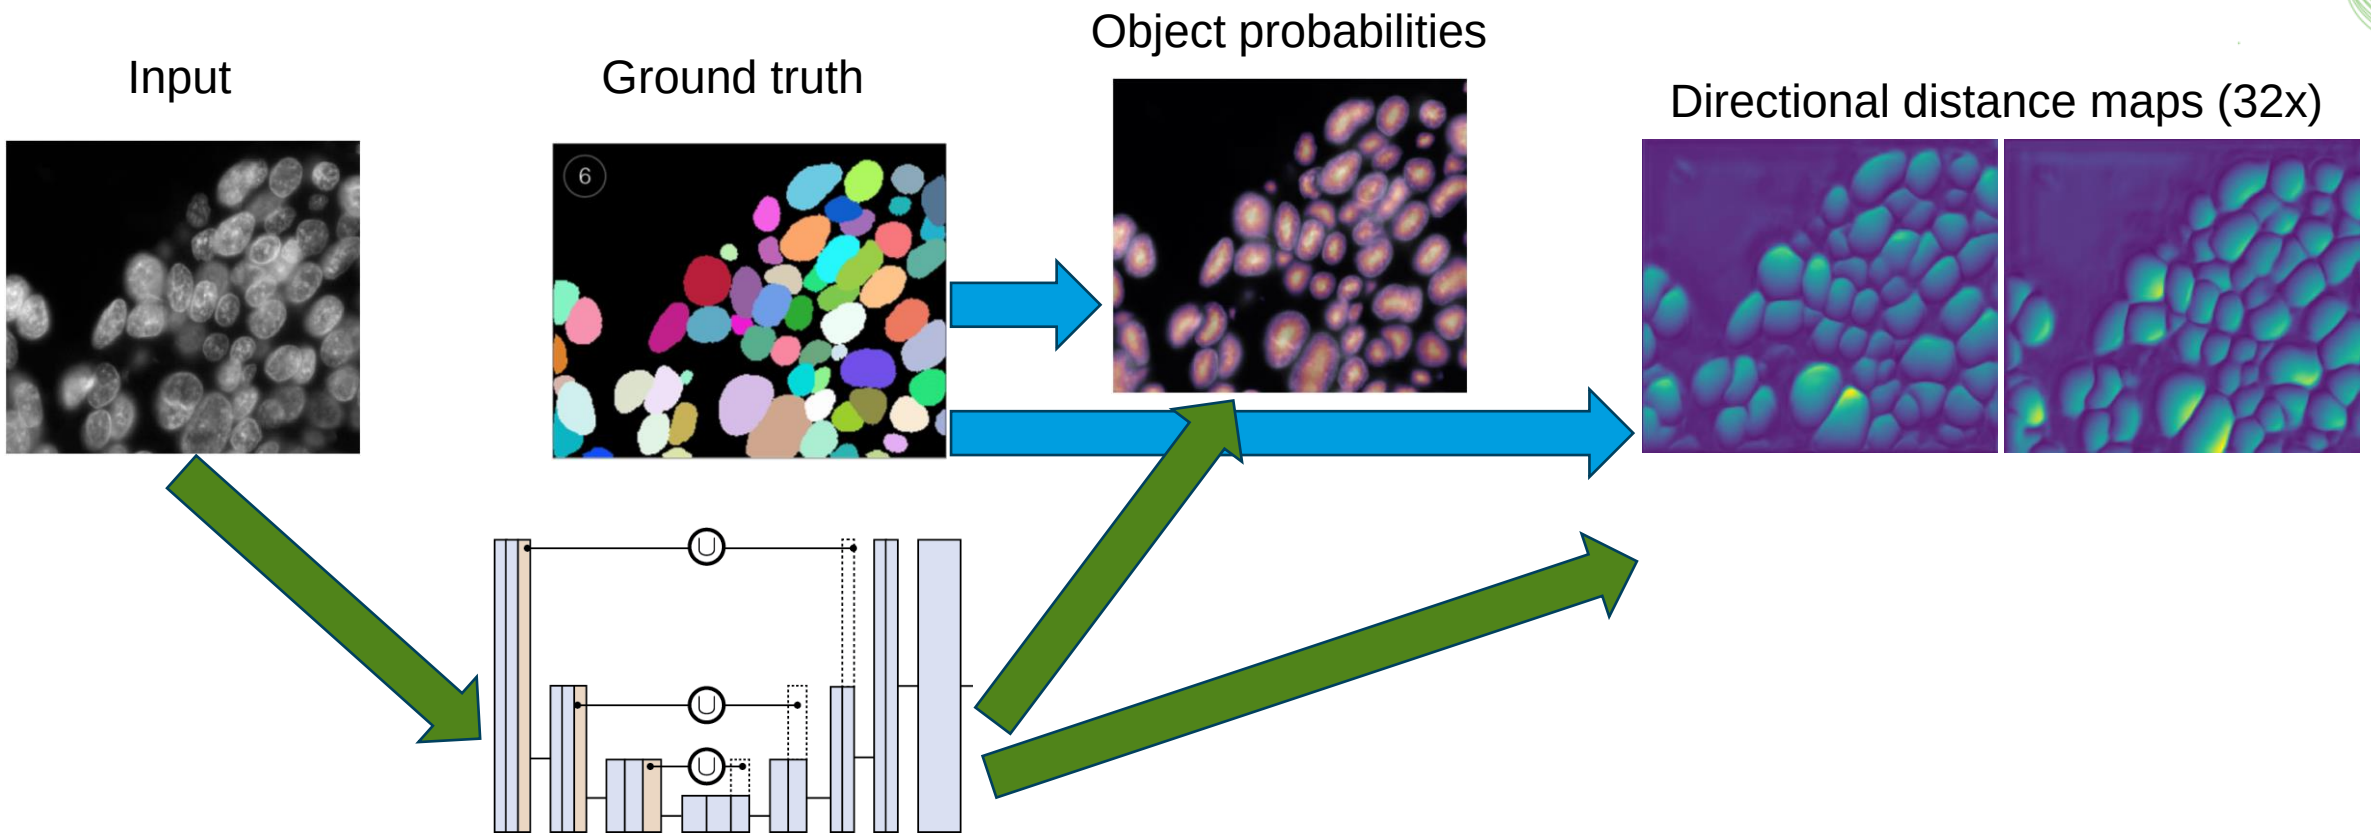

StarDist

- Prediction of probable object centers + polygon outlines
- Non-maximum-suppression of less likely polygons

Object probabilities

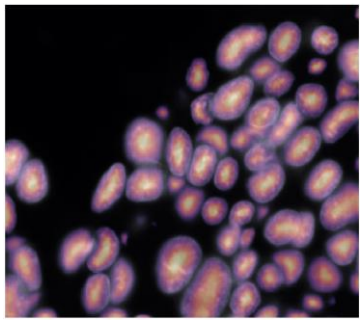

Directional distance maps

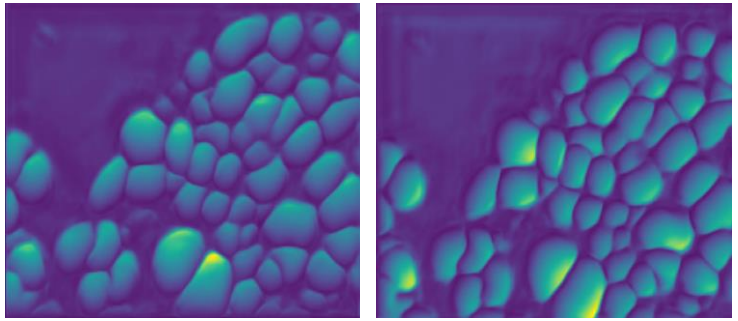

Polygon candidates

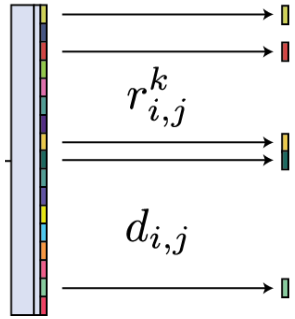

Final segmentation

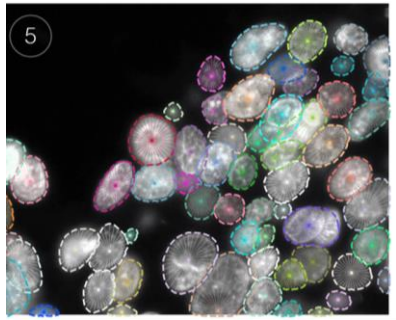

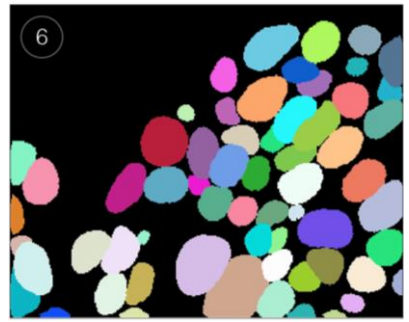

Slide adapted from: Martin Weigert, EPFL Lausanne and from https://github.com/stardist/stardist (license: BSD3)

StarDist: Limitations

Star-convex shapes!

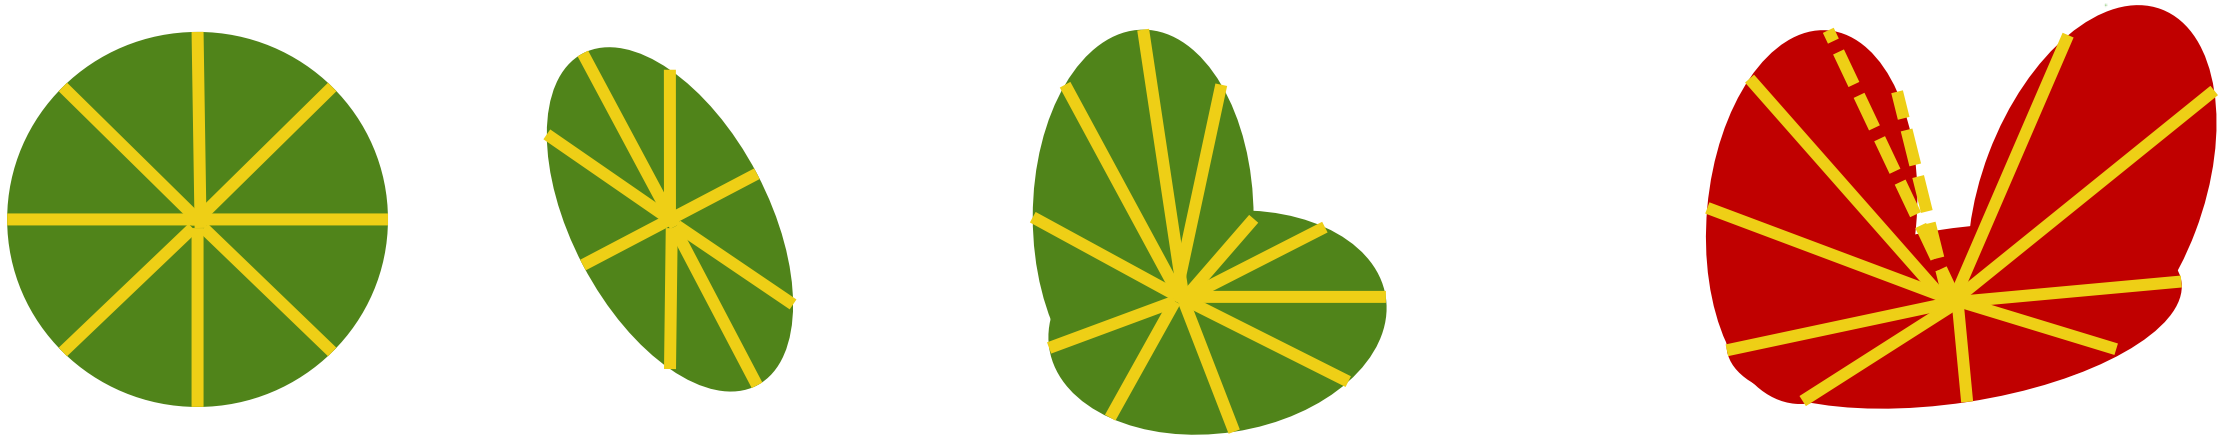

Good for nuclei, bad for cells.

StarDist: Python

Initialize model / download pretrained model

```
[3]: # creates a pretrained model
     model = StarDist2D.from_pretrained('2D_versatile_fluo')
```

```
Found model '2D_versatile_fluo' for 'StarDist2D'.
Loading network weights from 'weights_best.h5'.
Loading thresholds from 'thresholds.json'.
Using default values: prob_thresh=0.479071, nms_thresh=0.3.
```

Normalize intensity to range [0, 1]

```
[4]: axis_norm = (0,1)
     image = normalize(image, 1,99.8, axis=axis_norm)
```

Apply model

```
[5]: labels, details = model.predict_instances(image)

     stackview.insight(labels)
```

[5]:

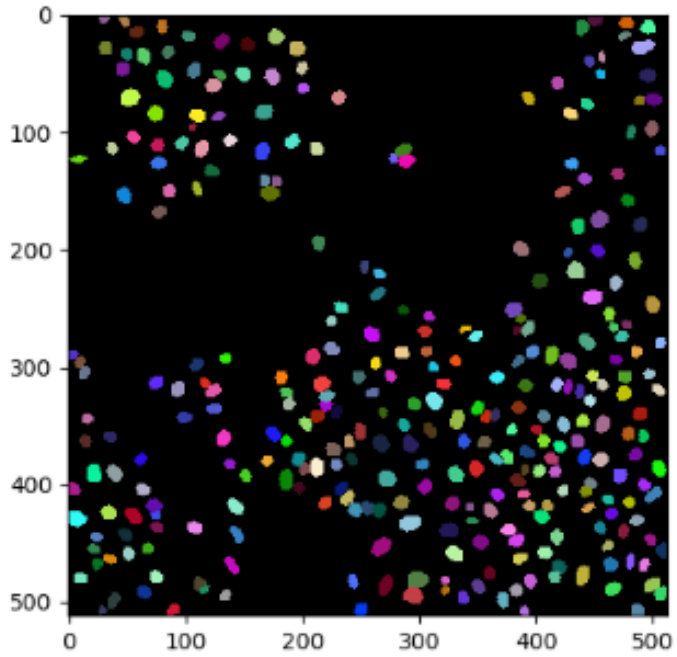

StarDist: Business model

Serving clinical sicentists using software based on open-source code

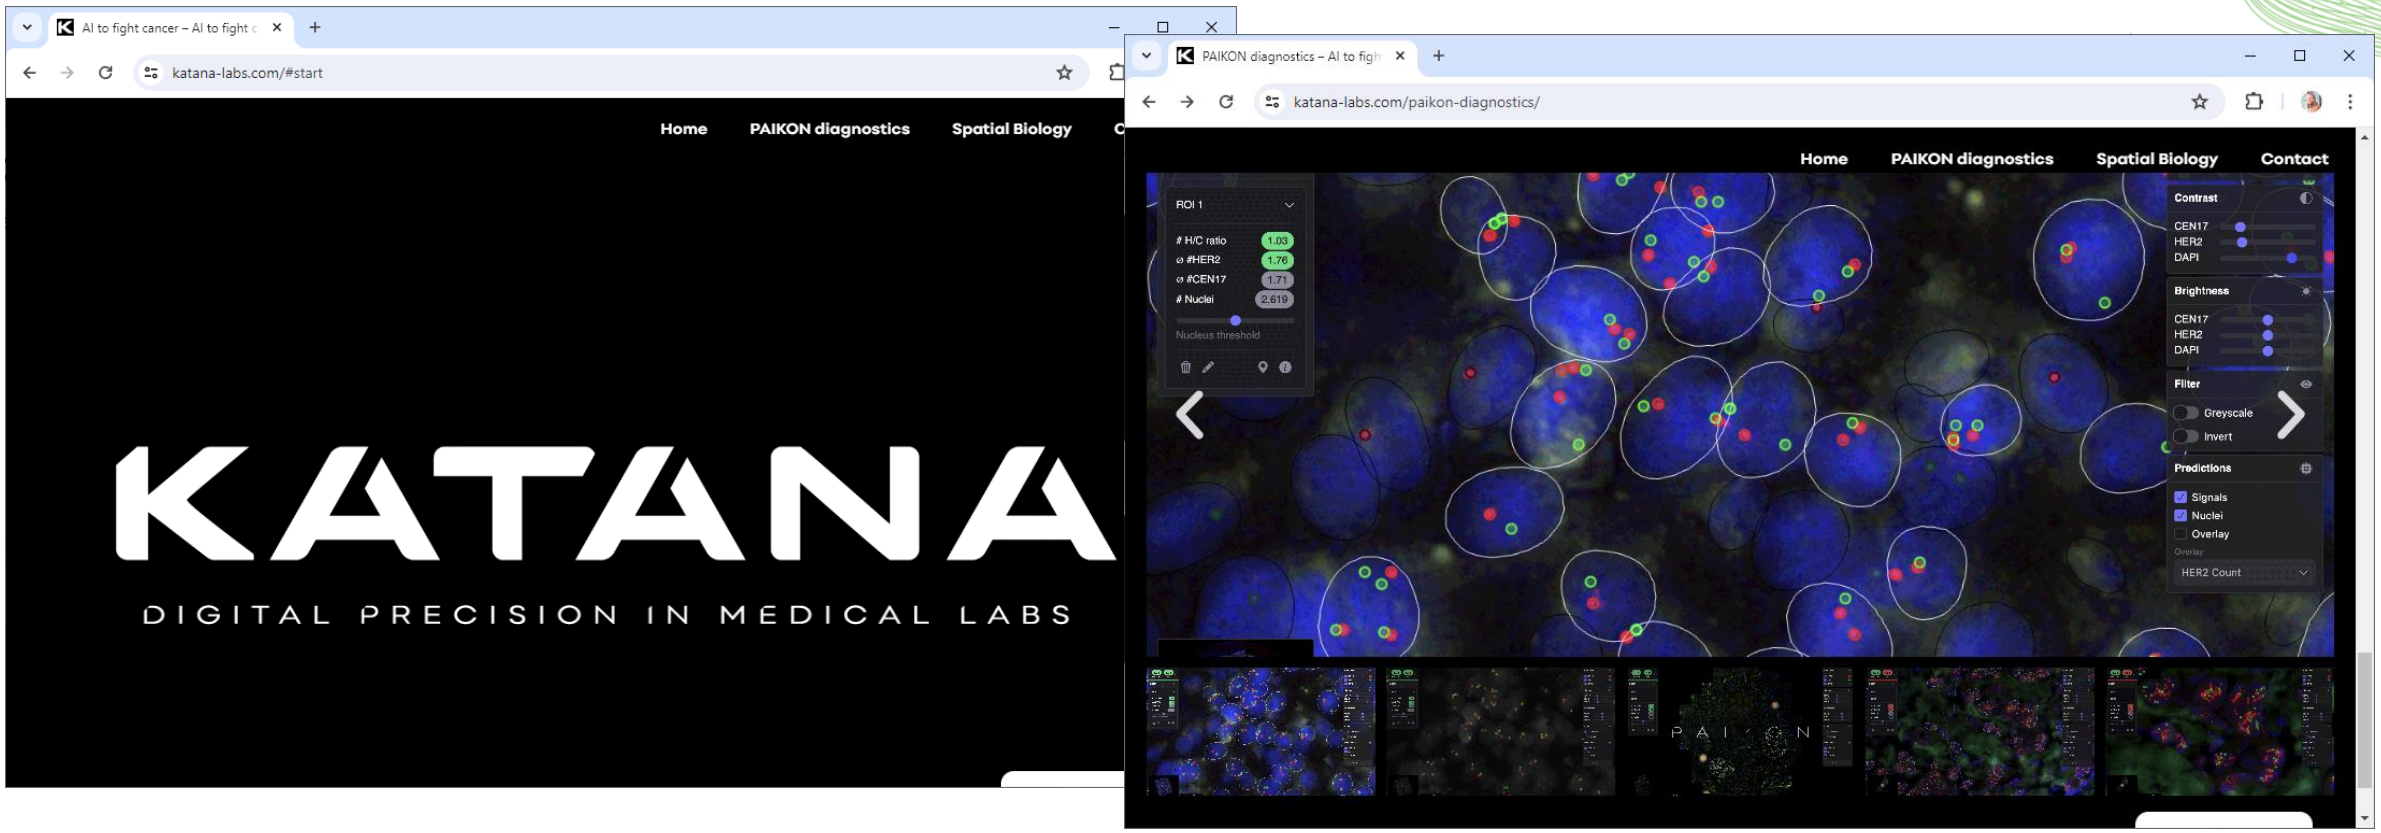

CellPose

- Processing more diverse datasets and object shapes

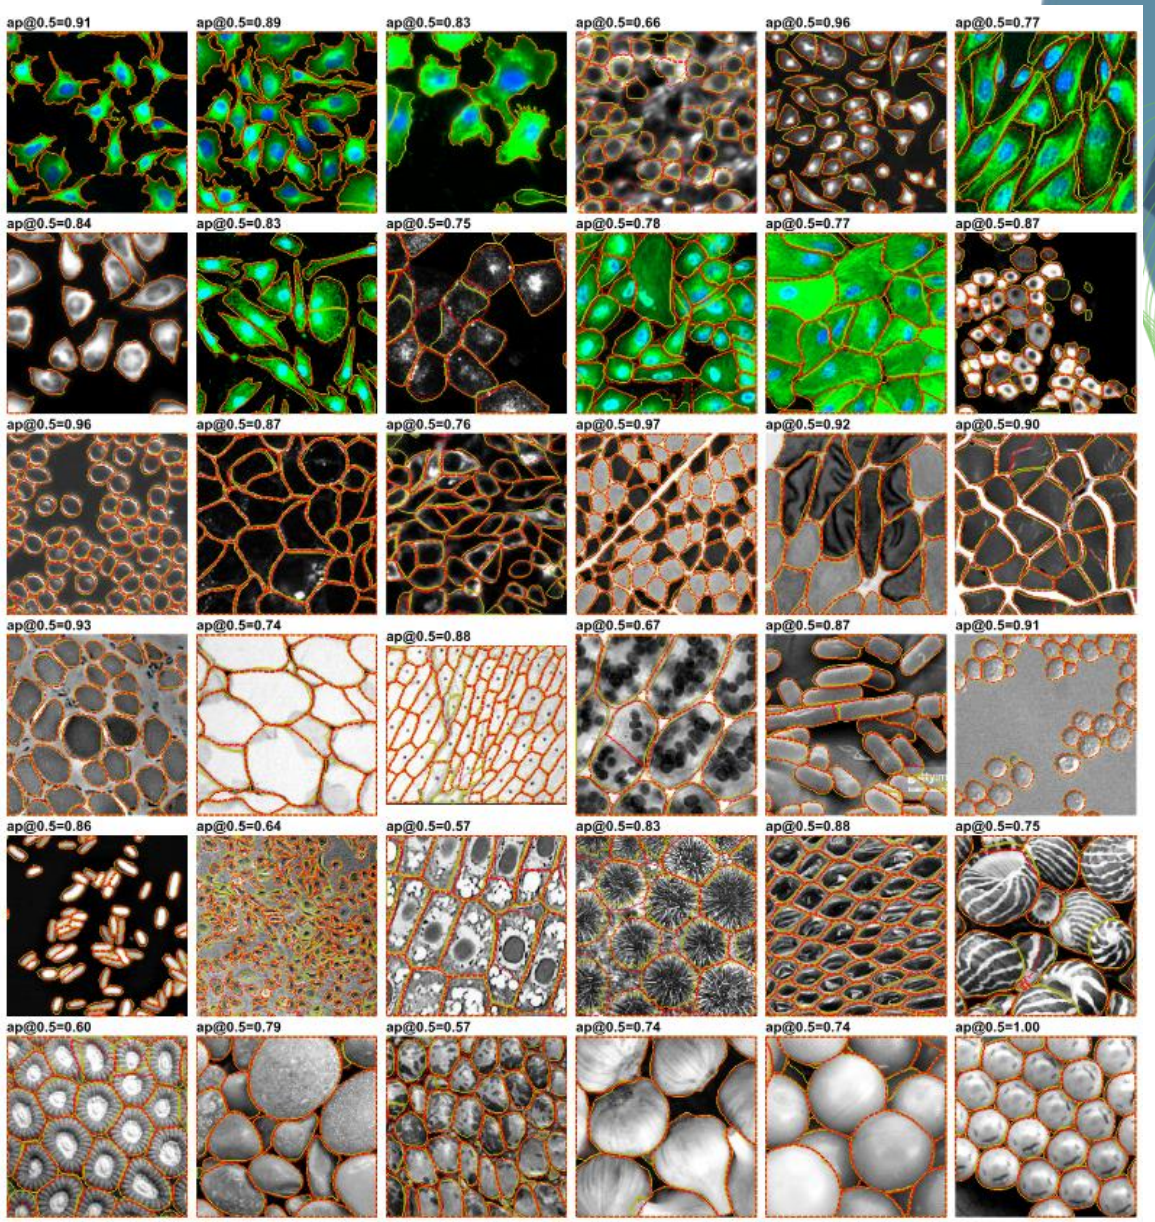

CellPose

Compute „flow fields“ from images

- using classical image processing during training
- using neural networks during prediction

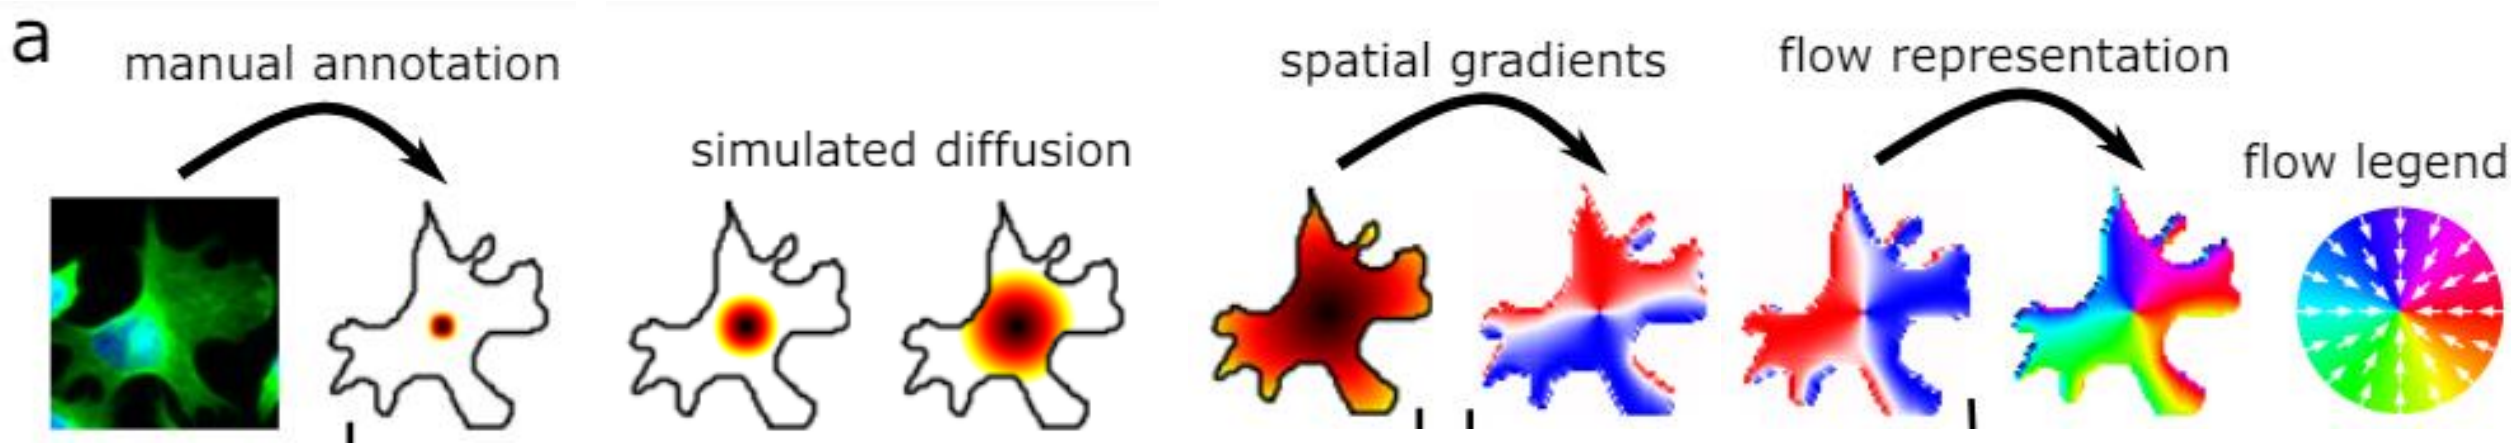

CellPose

- Cell/Nuclei – segmentation based on flow-fields
- Technically similar to Watershed, but with a deep-learning based altitude-image

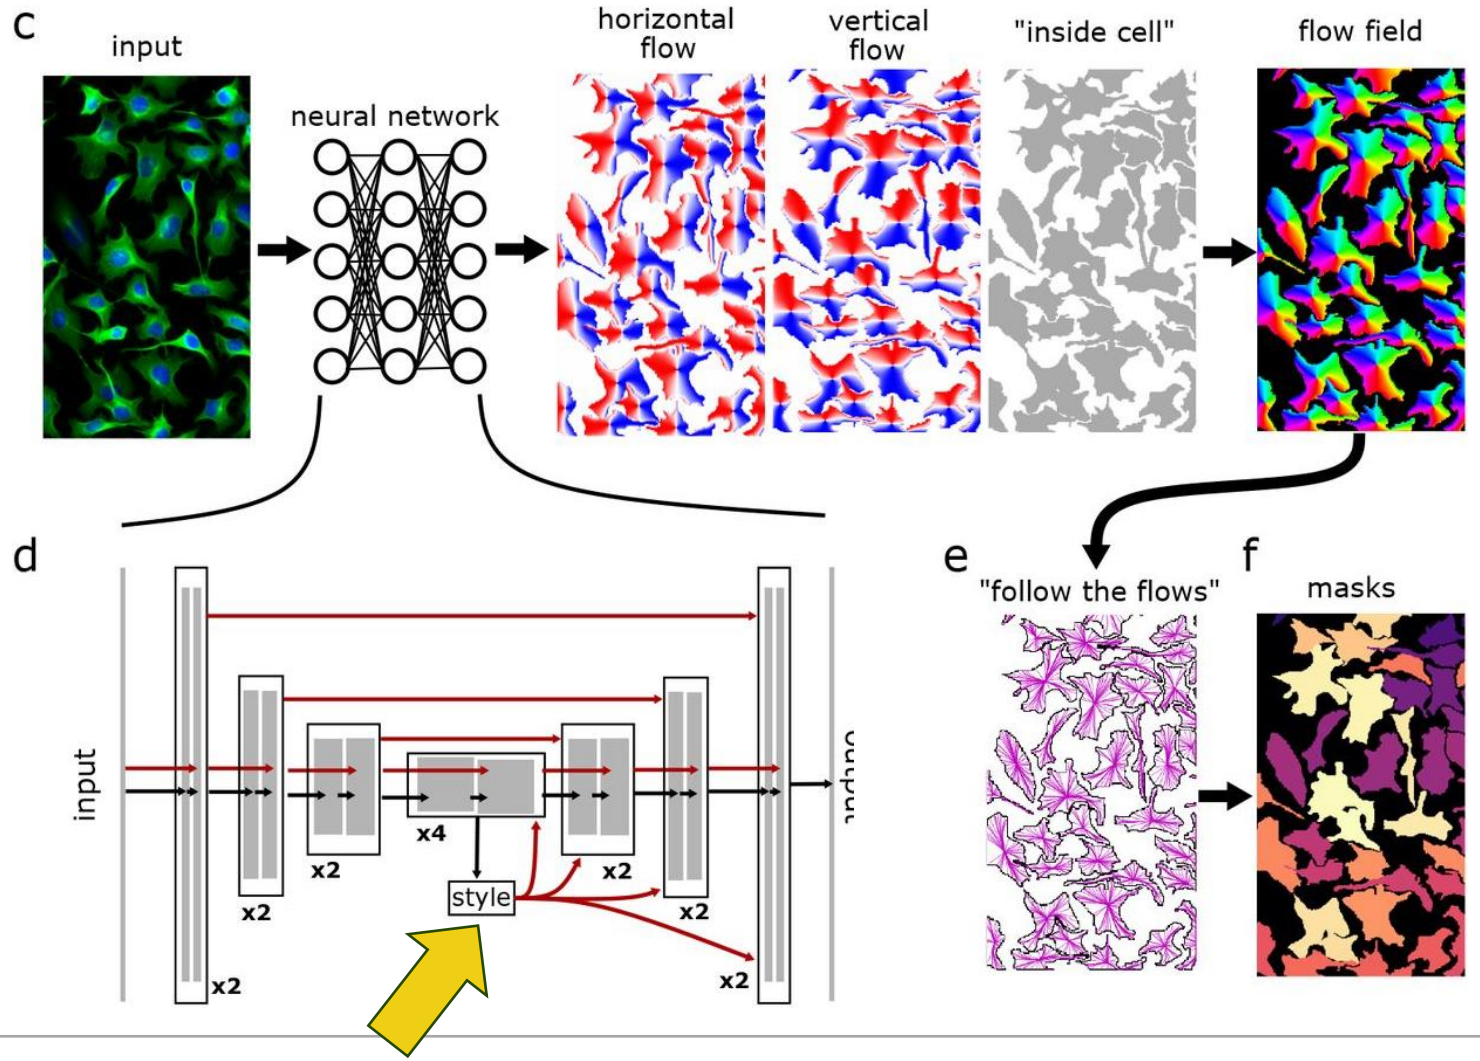

Image source: Modified from Stringer et al (2020) licensed CC-BY-NC 4.0
https://www.biorxiv.org/content/10.1101/2020.02.02.931238v2.full

CellPose

- Image *style* is a parameter determined before prediction to guide segmentation.

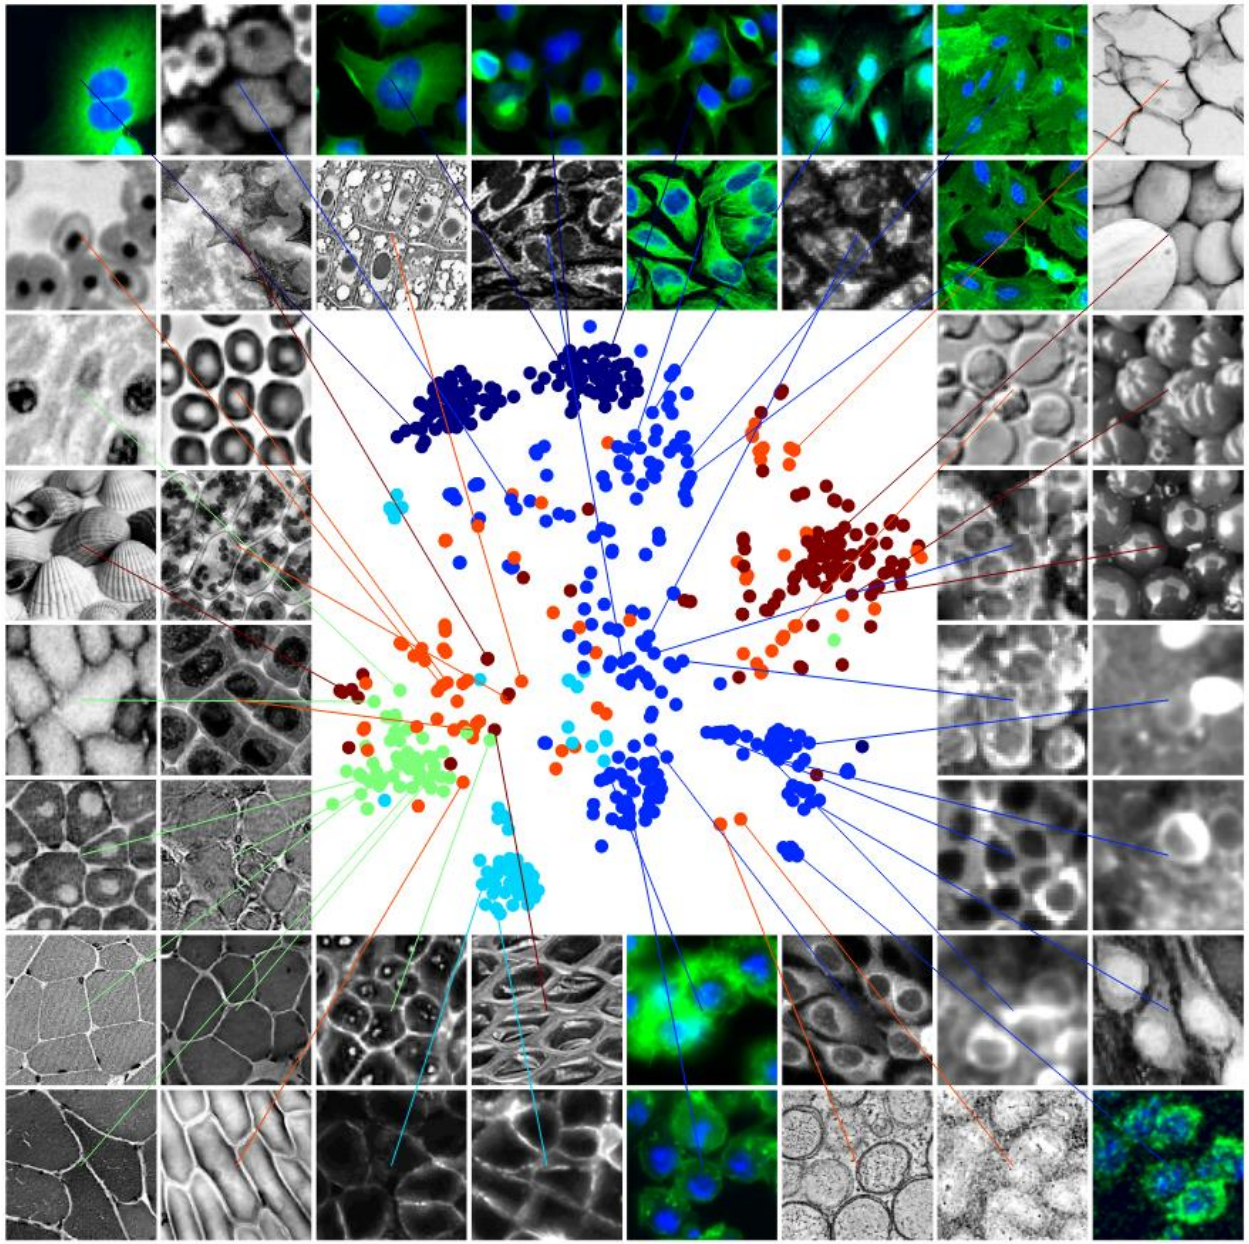

Cellpose 3D

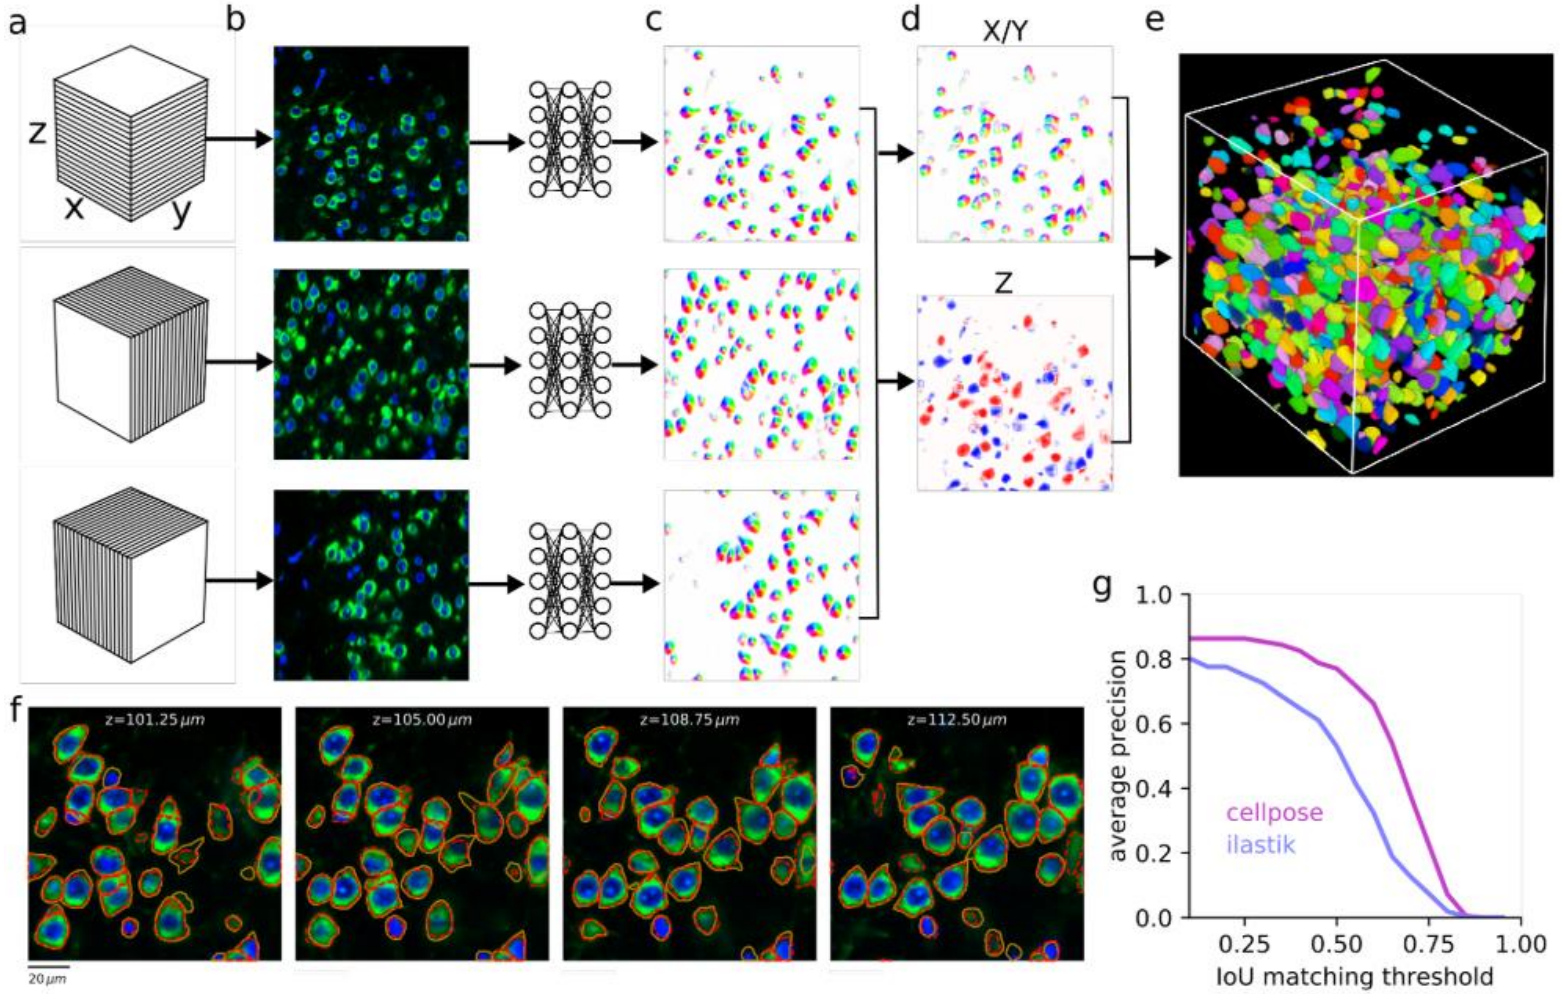

Cellpose in Napari

mamba install cellpose-napari

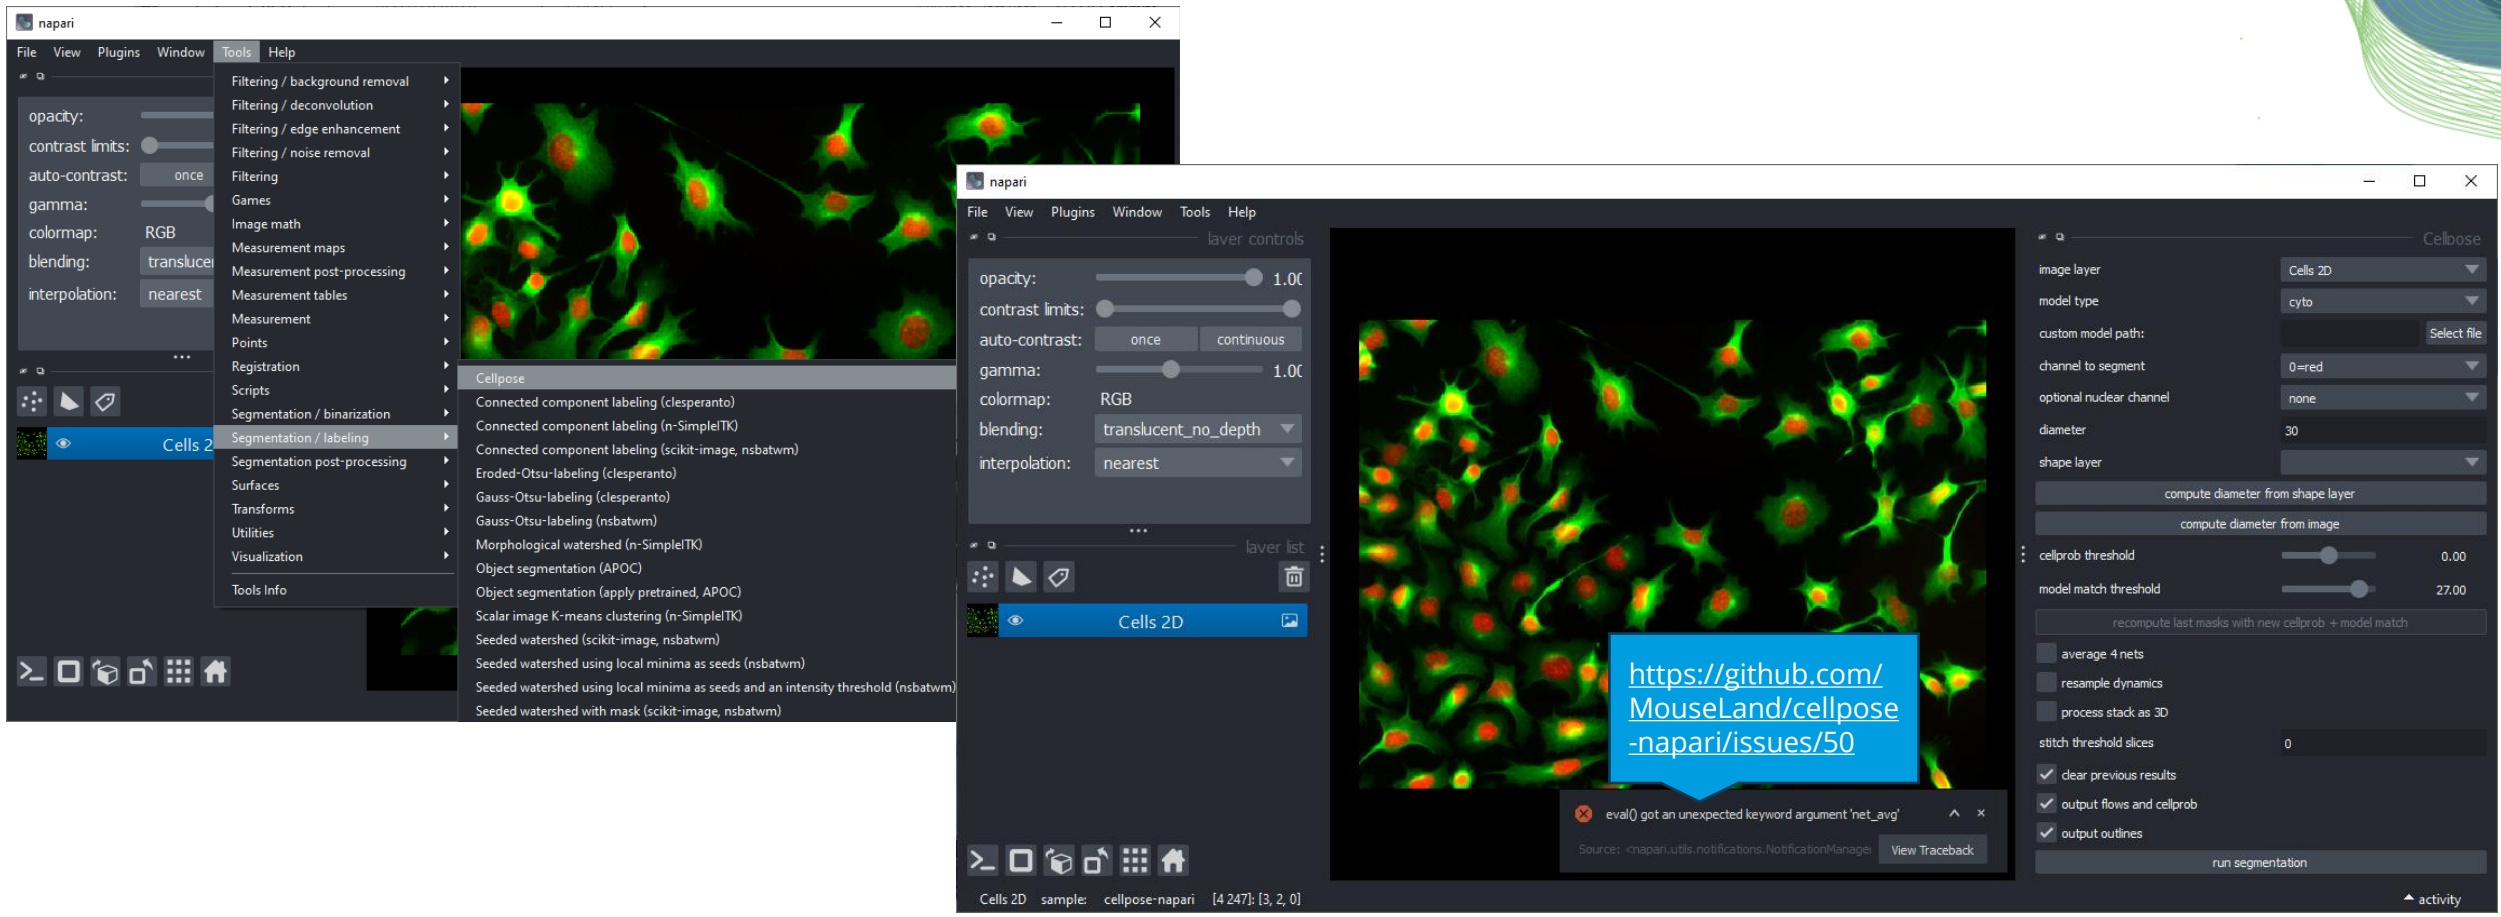

Cellpose in Python

Initialize a pretrained model

```
[3]: model = cellpose.models.Cellpose(gpu=False, model_type='nuclei')
```

List available models

```
[8]: cellpose.models.MODEL_NAMES
```

```
[8]: ['cyto3',
      'nuclei',
      'cyto2_cp3',
      'tissuenet_cp3',
      'livecell_cp3',
      'yeast_PhC_cp3',
      'yeast_BF_cp3',
      'bact_phase_cp3',
      'bact_fluor_cp3',
      'deepbacs_cp3',
      'cyto2',
      'cyto',
      'transformer_cp3',
      'neurips_cellpose_default',
      'neurips_cellpose_transformer',
      'neurips_grayscale_cyto2']
```

Cellpose in Python

```
[4]: channels = [0,0] # This means we are processing single-channel greyscale images.

masks, flows, styles, diams = model.eval(image, diameter=None, channels=channels)
```

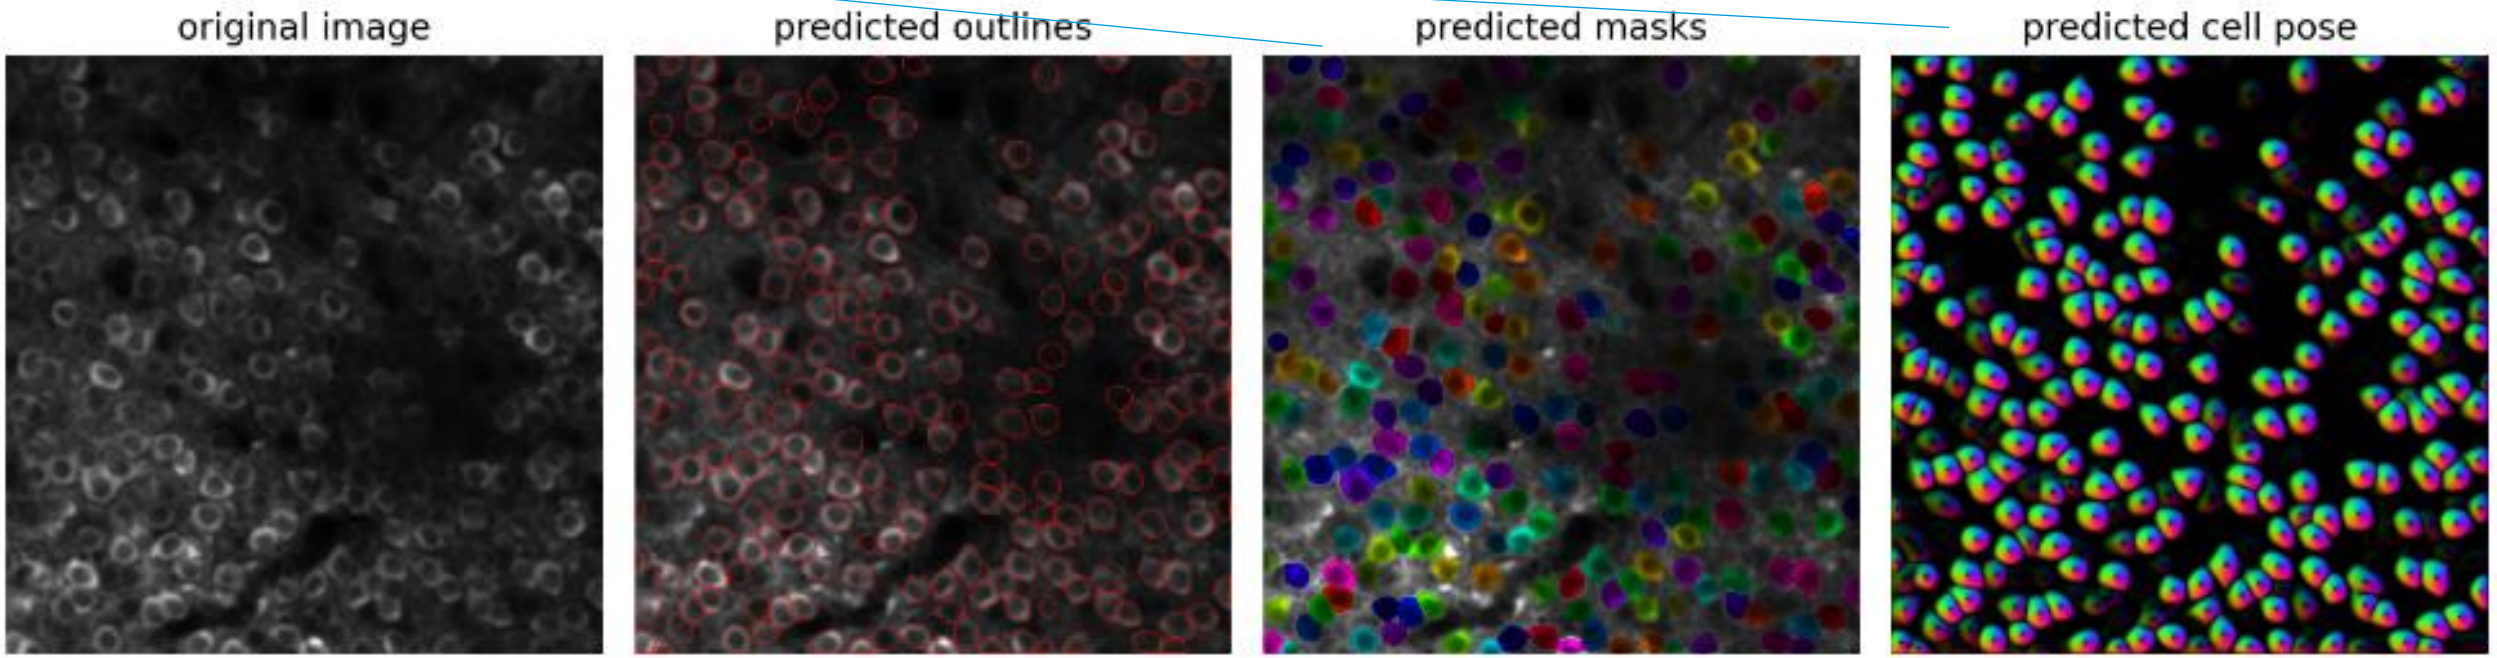

Cellpose 2

Train your own model

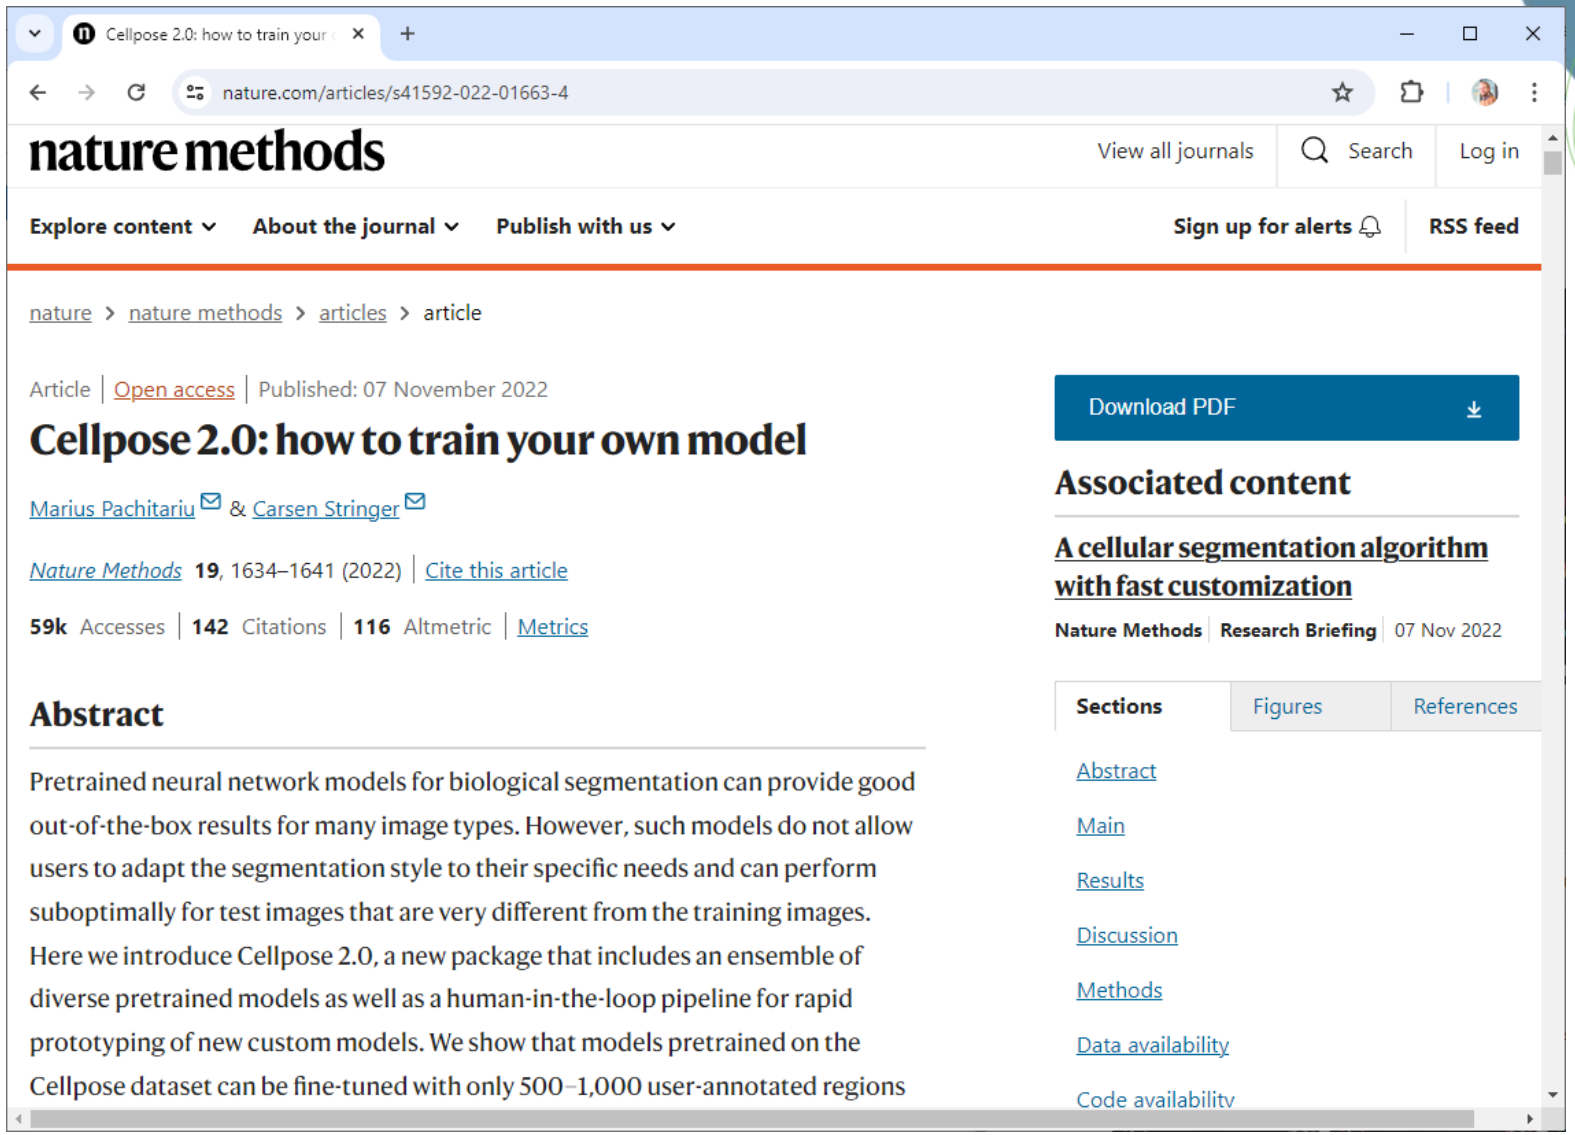

Cellpose 3

Cellpose 3 includes denoising and deblurring, to improve image segmenation quality

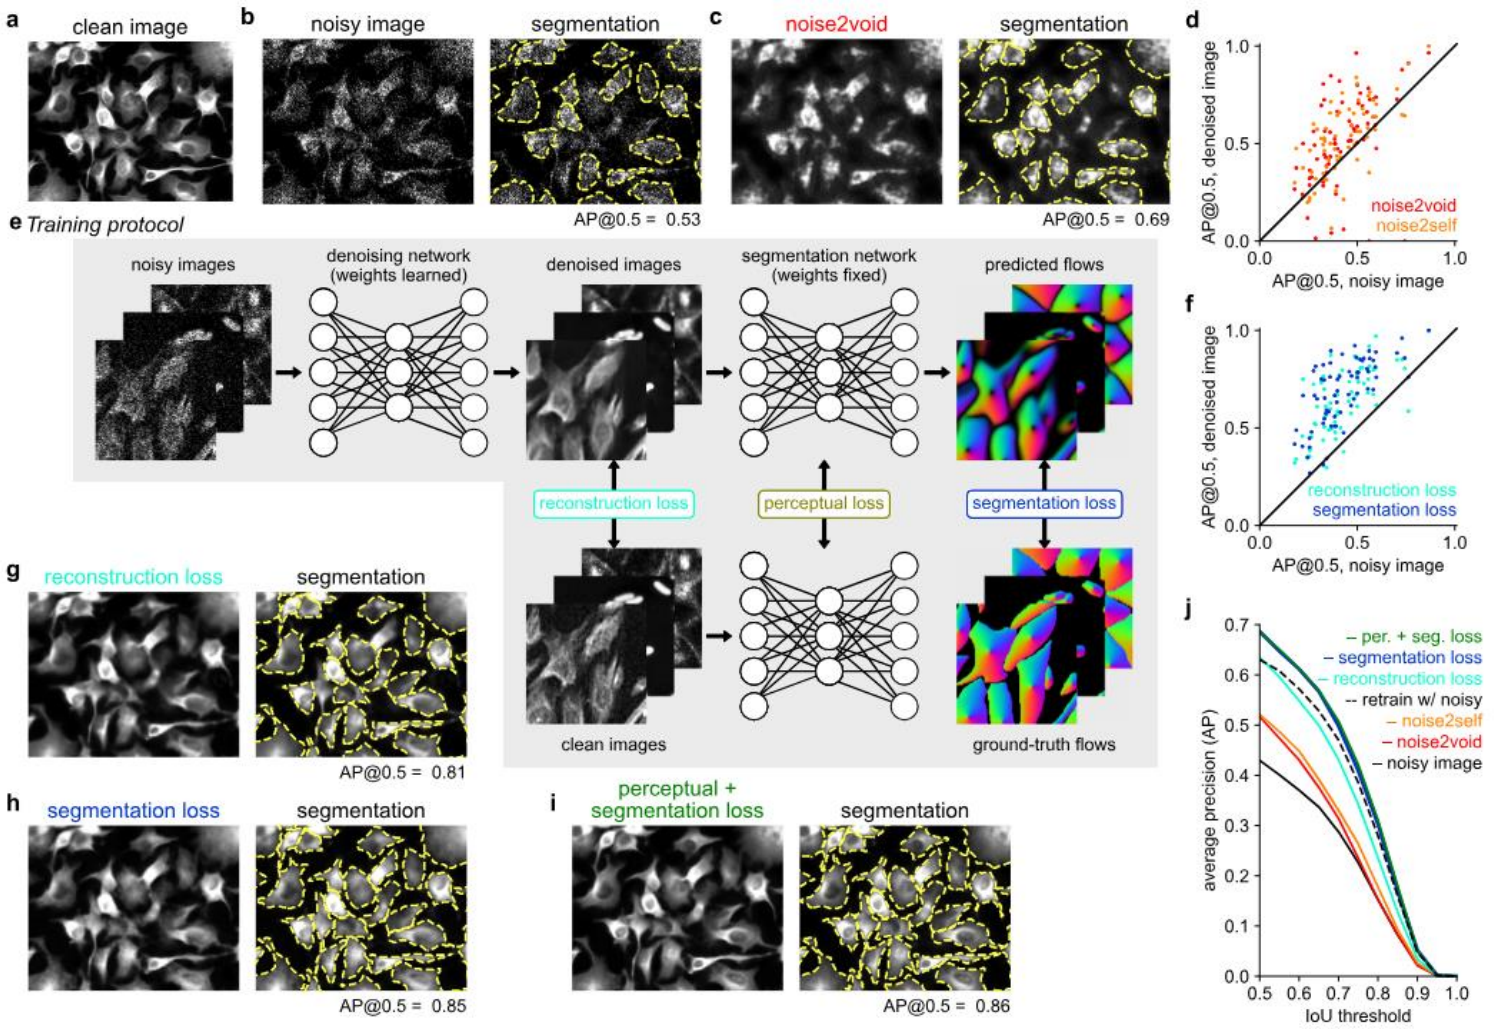

Cellpose

State-of-the-art despite the rise of transformers

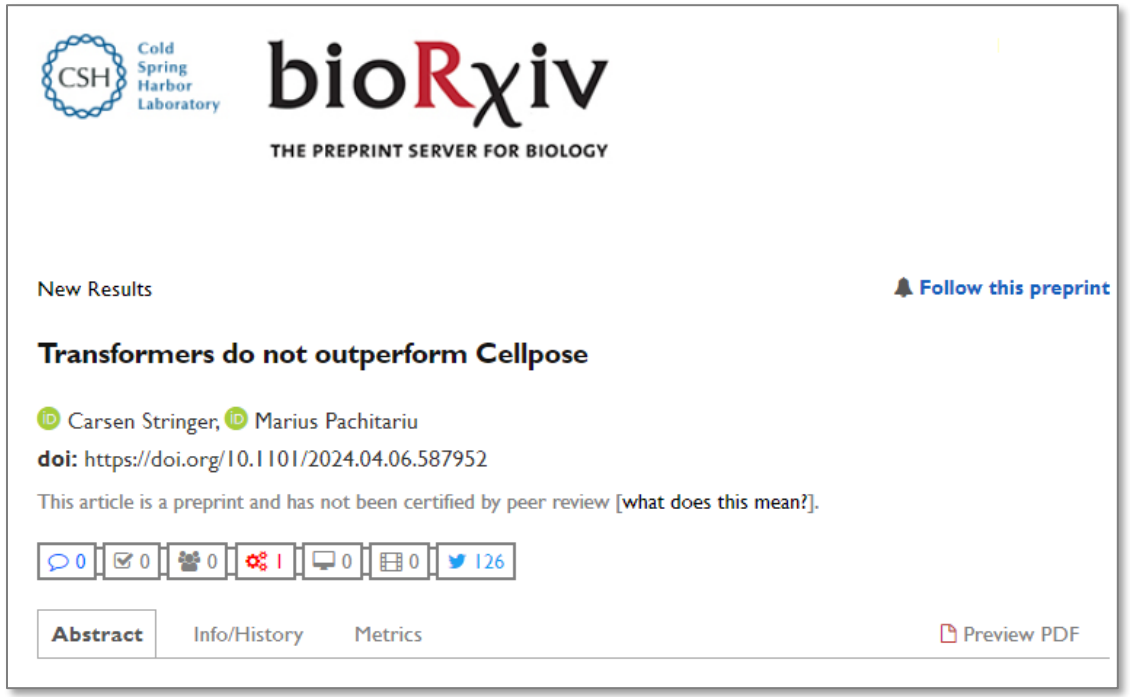

Cold Spring Harbor Laboratory

bioRxiv

THE PREPRINT SERVER FOR BIOLOGY

New Results

Follow this preprint

Transformers do not outperform Cellpose

Carsen Stringer, Marius Pachitariu

doi: https://doi.org/10.1101/2024.04.06.587952

This article is a preprint and has not been certified by peer review [what does this mean?].

0 0 0 1 0 0 126

Abstract Info/History Metrics Preview PDF

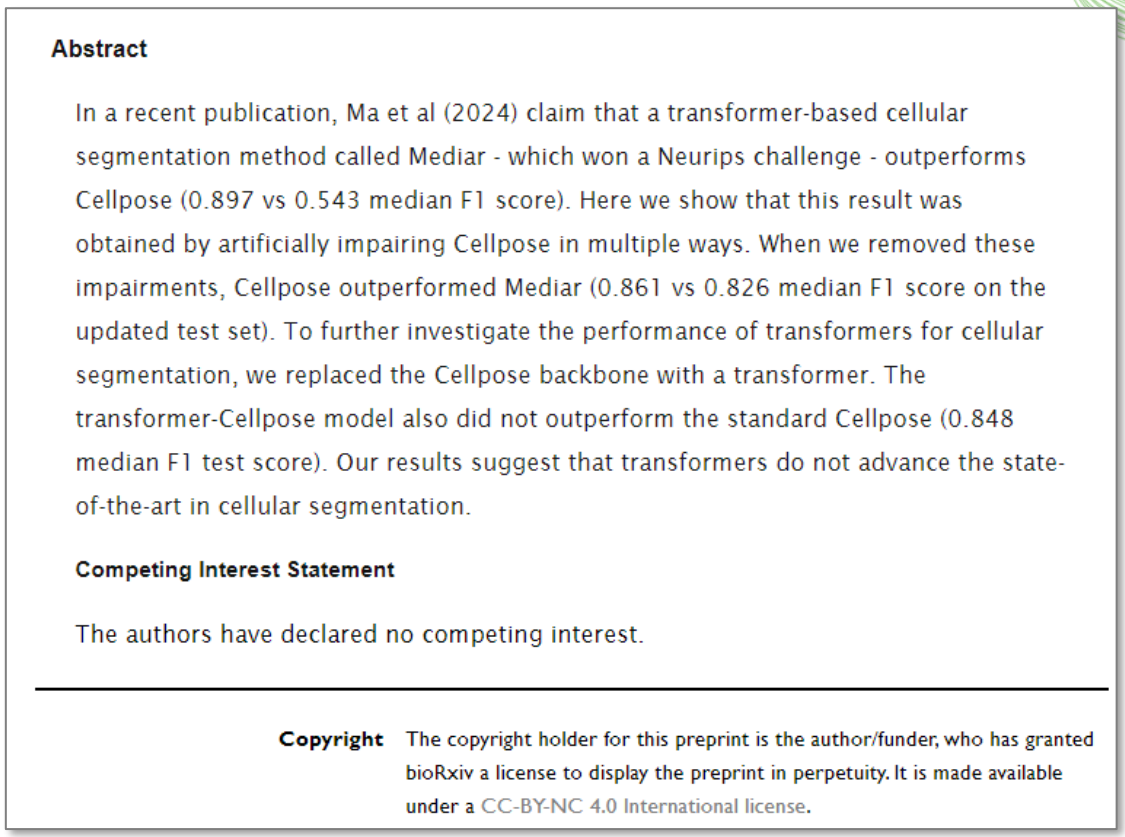

Abstract

In a recent publication, Ma et al (2024) claim that a transformer-based cellular segmentation method called Mediar - which won a Neurips challenge - outperforms Cellpose (0.897 vs 0.543 median F1 score). Here we show that this result was obtained by artificially impairing Cellpose in multiple ways. When we removed these impairments, Cellpose outperformed Mediar (0.861 vs 0.826 median F1 score on the updated test set). To further investigate the performance of transformers for cellular segmentation, we replaced the Cellpose backbone with a transformer. The transformer-Cellpose model also did not outperform the standard Cellpose (0.848 median F1 test score). Our results suggest that transformers do not advance the state-of-the-art in cellular segmentation.

Competing Interest Statement

The authors have declared no competing interest.

Segment Anything Model

New approach to DL-based image segmentation involving prompts

Segment Anything Model

New approach to DL-based image segmentation involving prompts

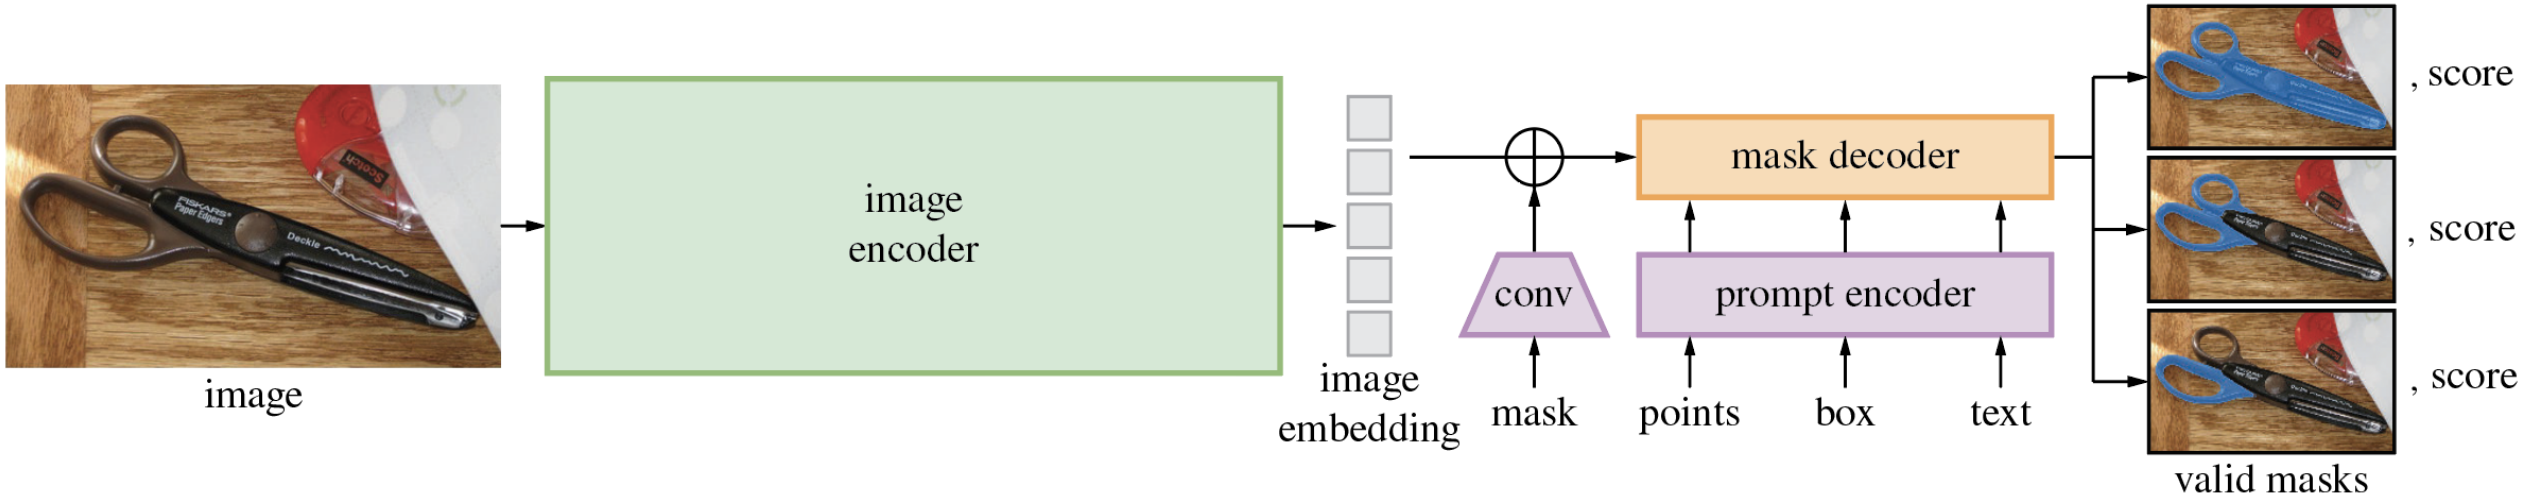

Source: https://github.com/facebookresearch/segment-anything

Segment Anything Model

New approach to DL-based image segmentation involving prompts

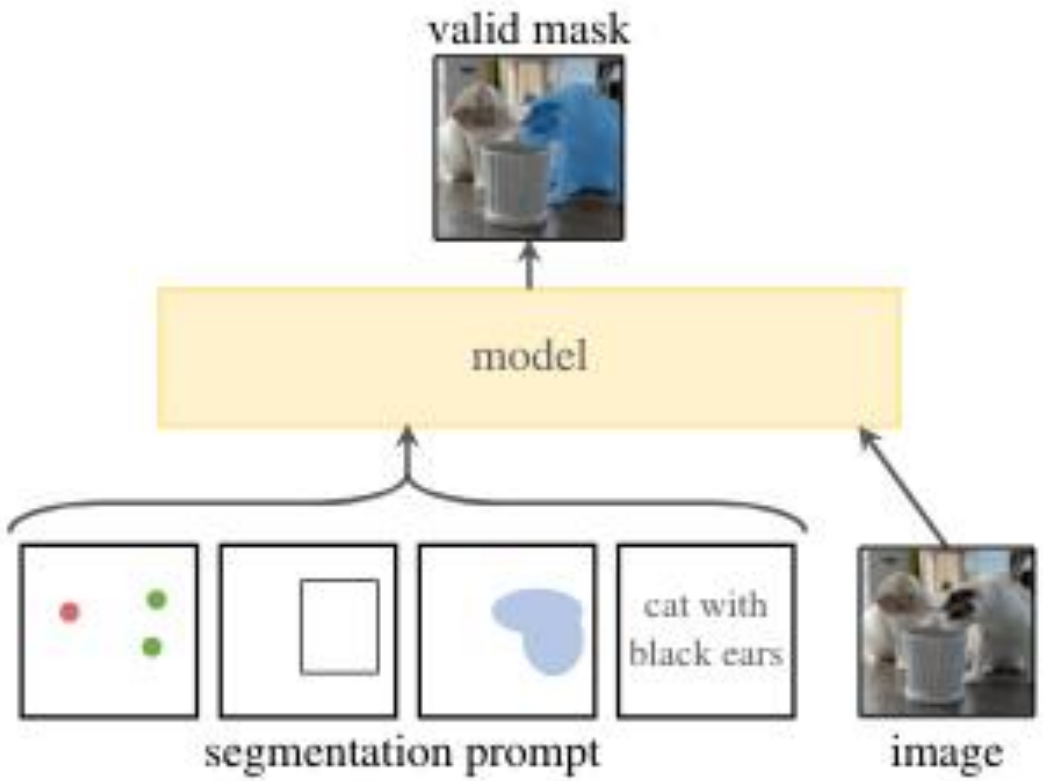

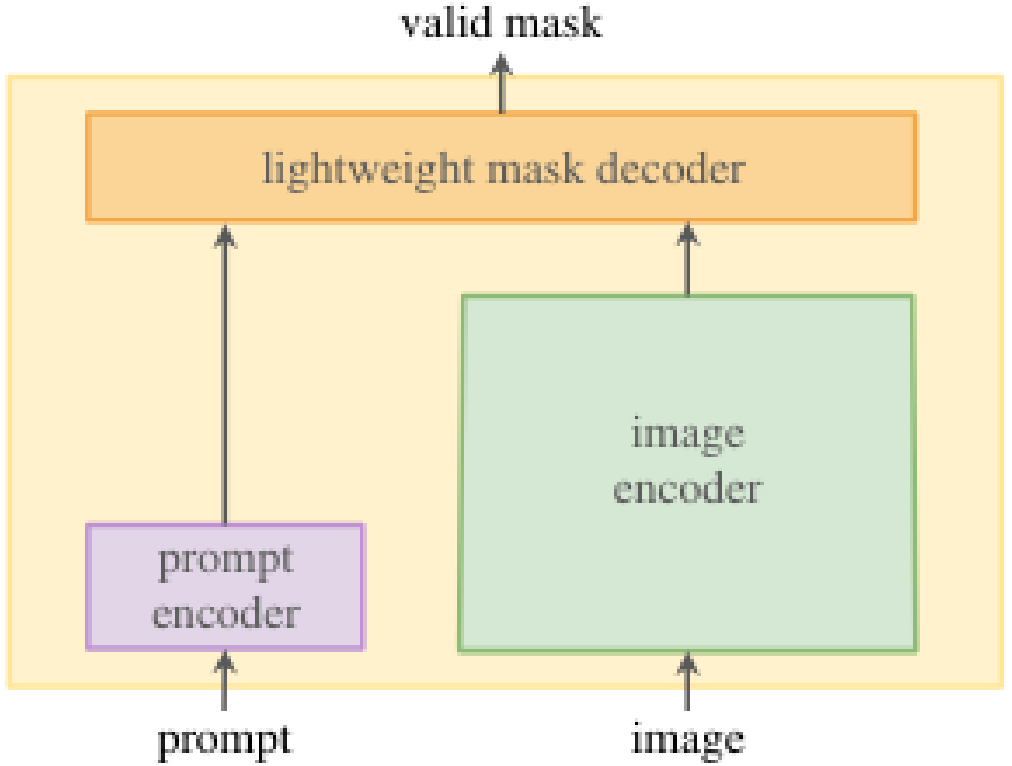

Segment Anything Model

Trained on mostly natural images

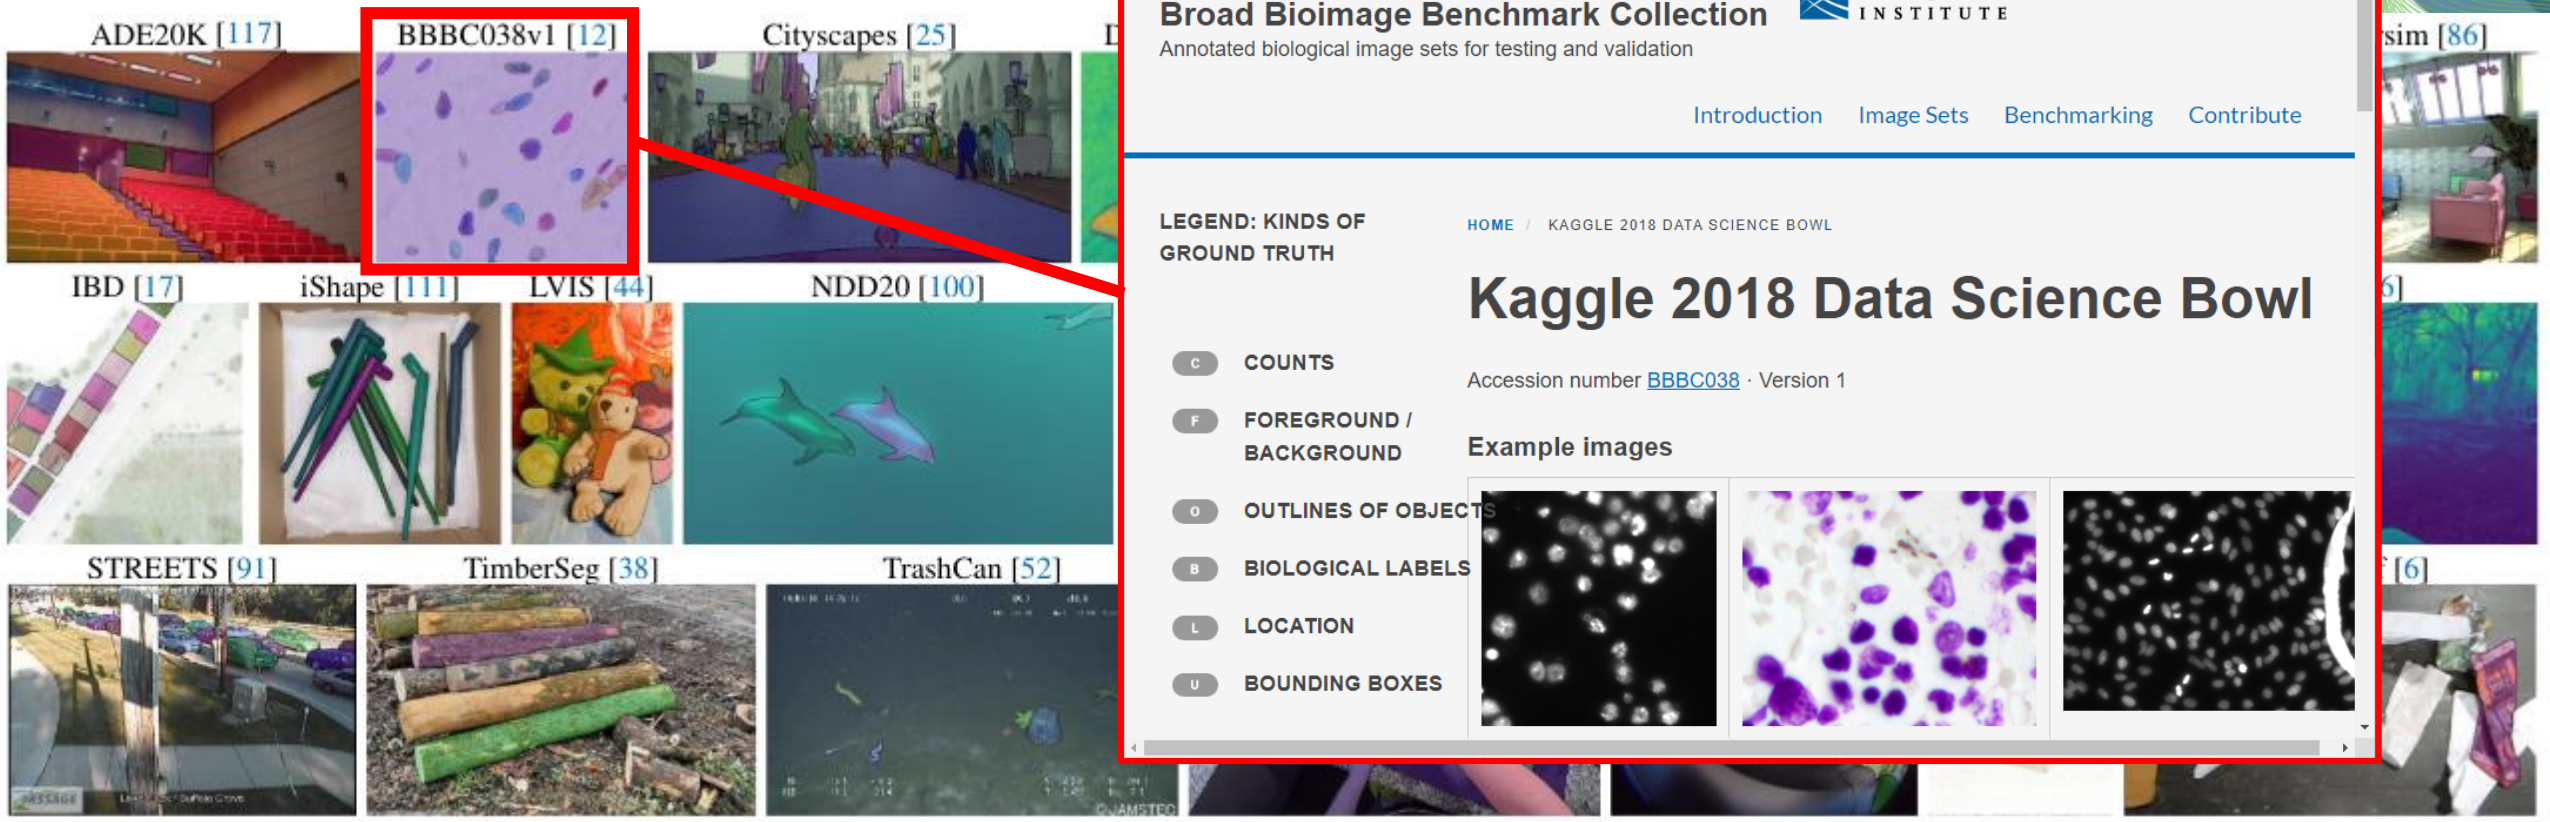

Figure 8: Samples from the 23 diverse segmentation datasets used to evaluate SAM's zero-shot transfer capabilities.

Segment Anything for Microscopy

Popping up napari plugins, some within 24h after SAM was published

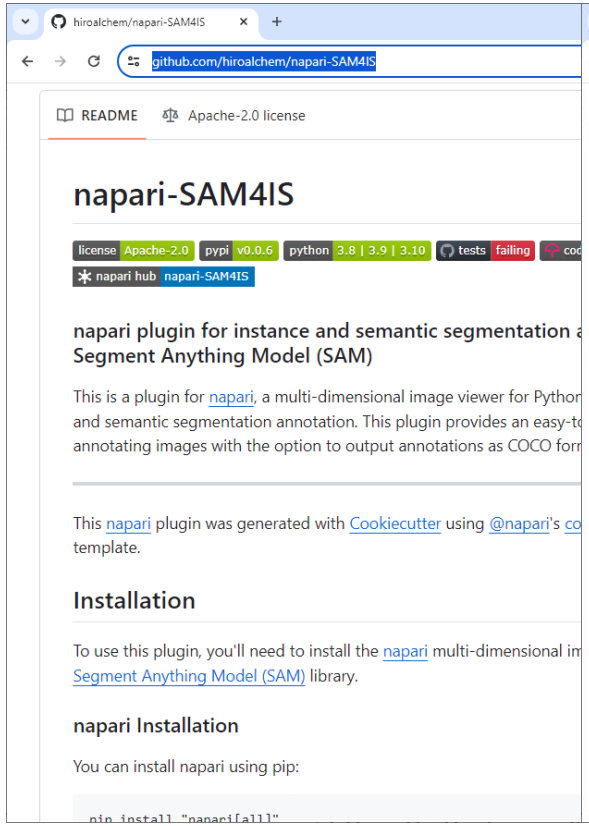

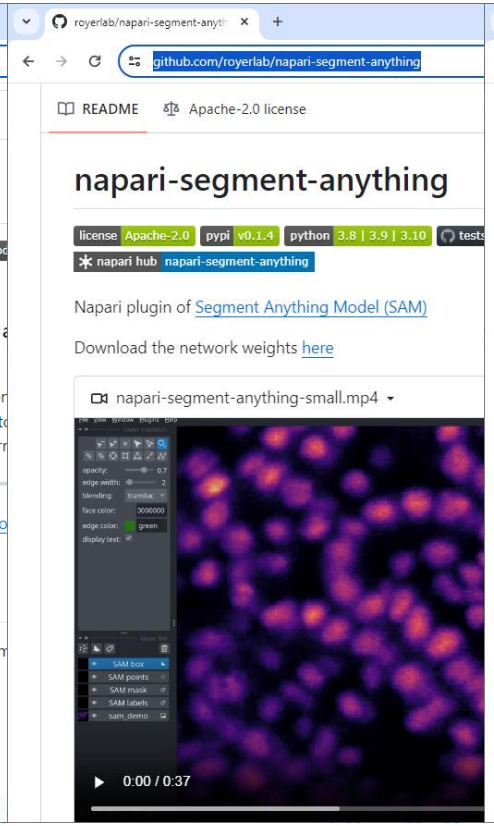

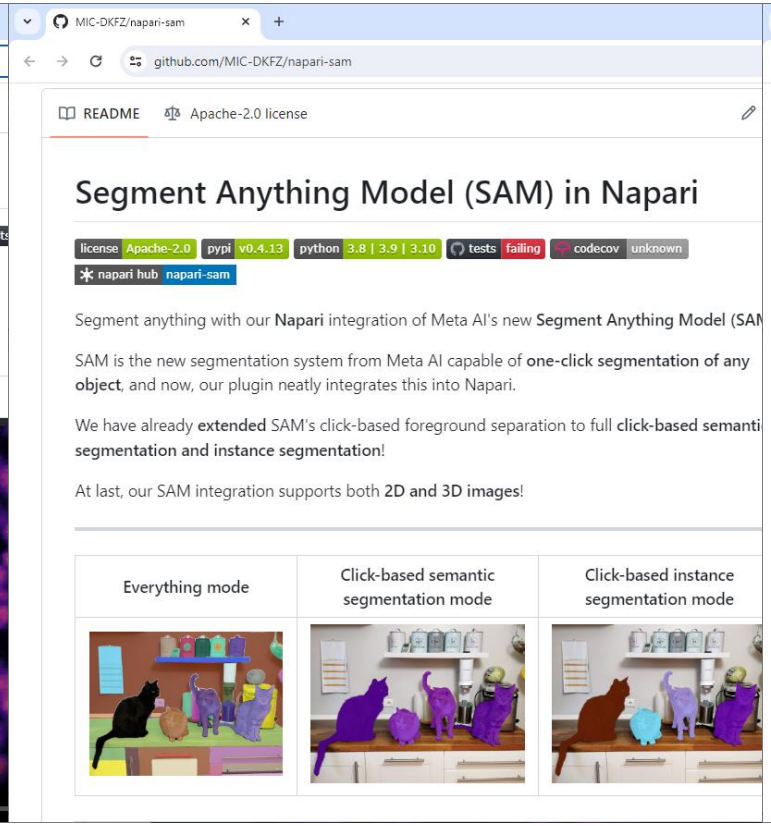

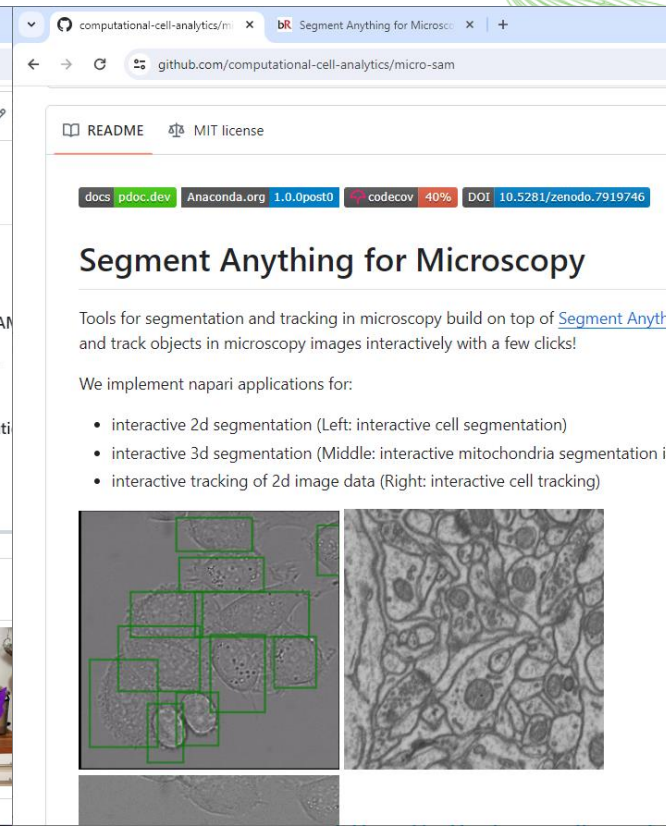

Segment Anything for Microscopy

Downsides:

- Original code did not contain the procedure for iterative training
- Instance segmentation not ideal (watershed-implementation added in micro-sam)
- Fine-tuning for microscopy data necessary

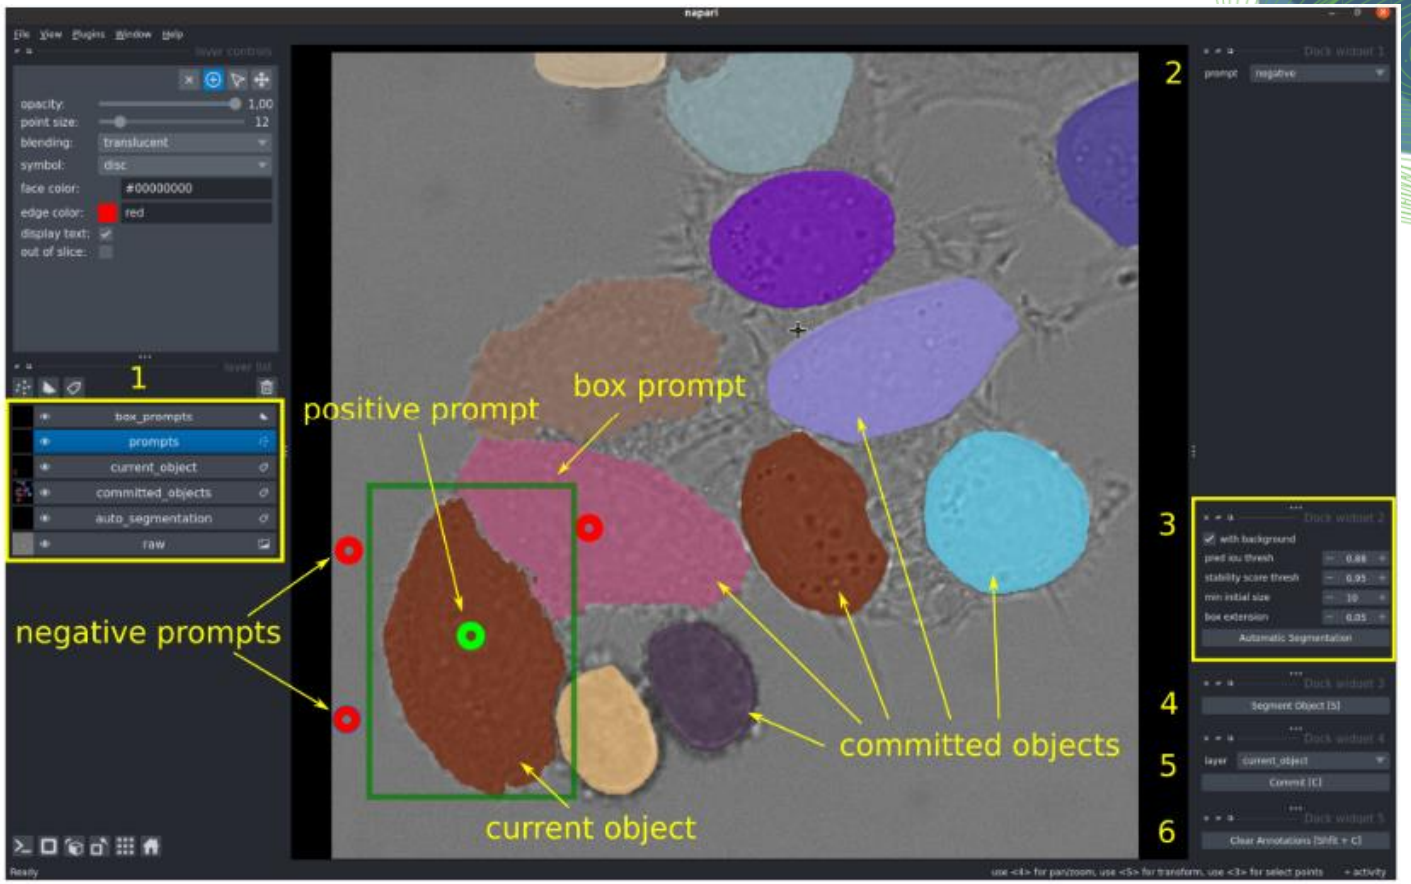

https://www.biorxiv.org/content/10.1101/2023.08.21.554208v1.article-info
https://github.com/computational-cell-analytics/micro-sam

Segment Anything for Microscopy

Real-world scenarios: human-in-the-loop

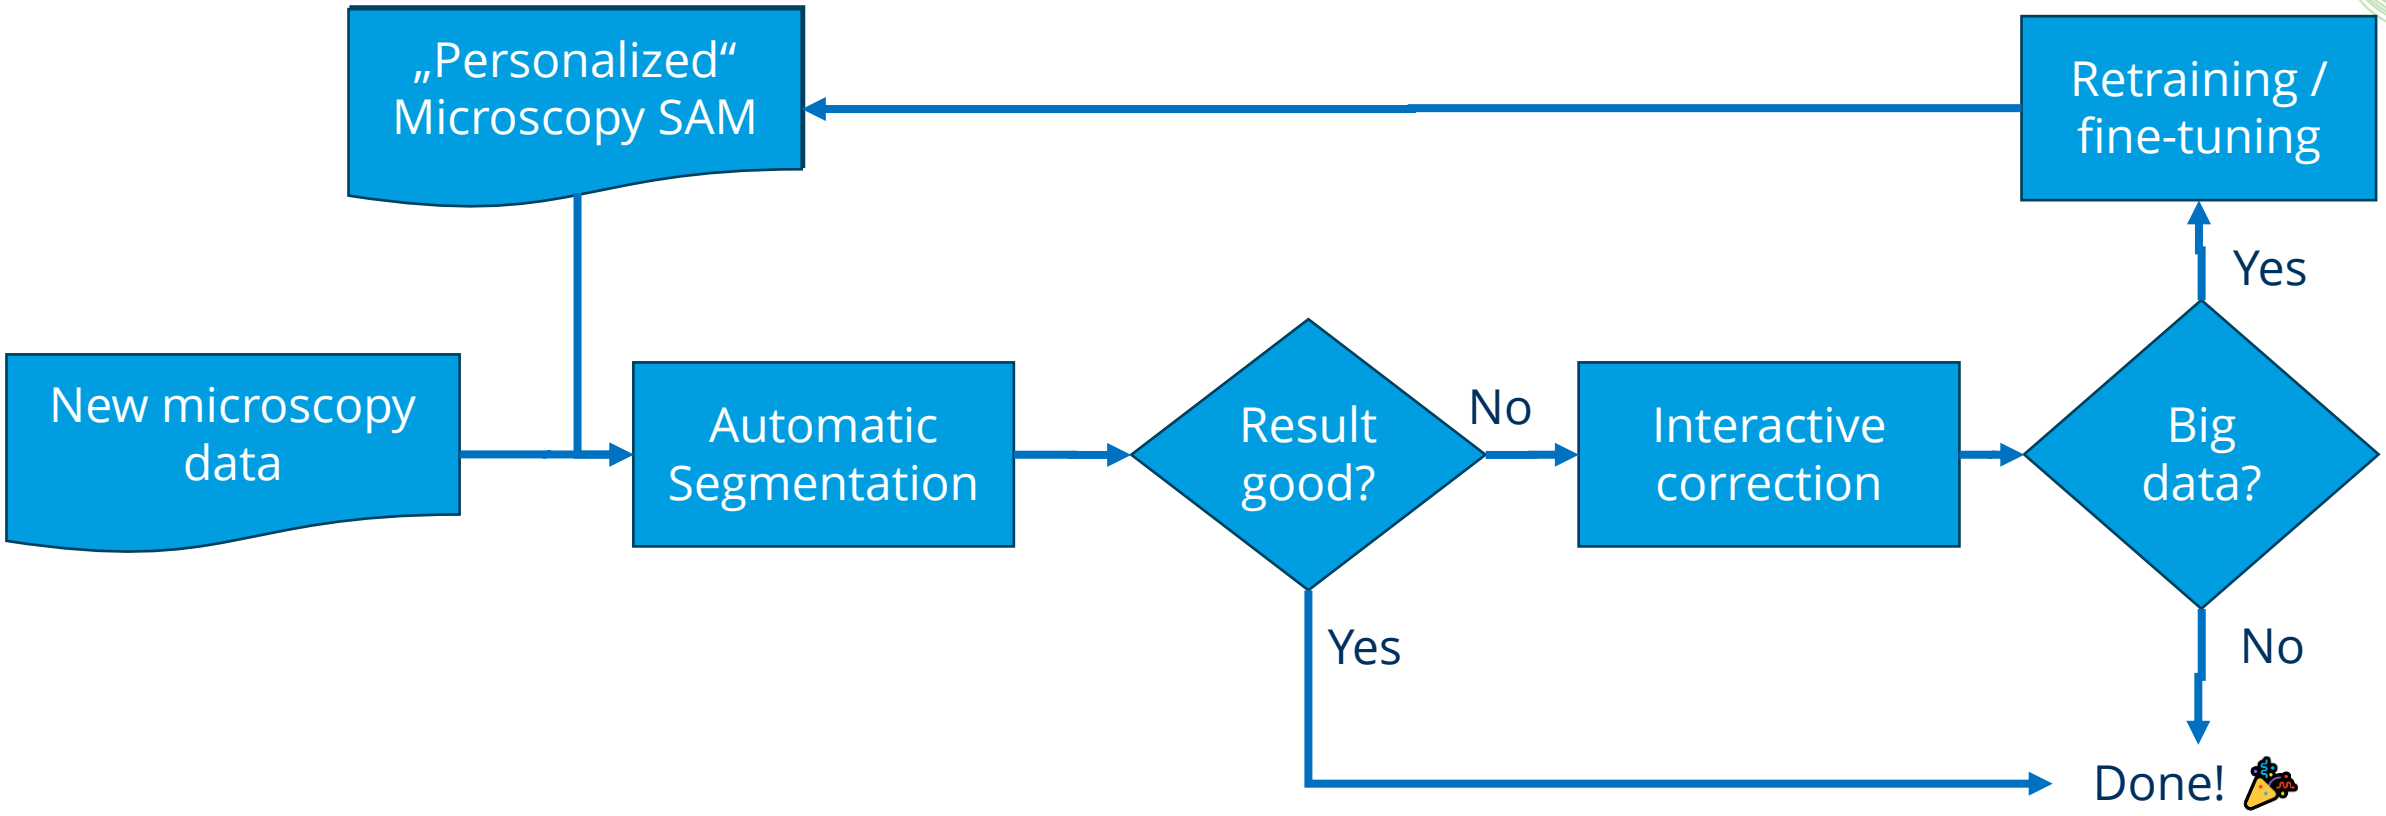

Segment Anything for Microscopy

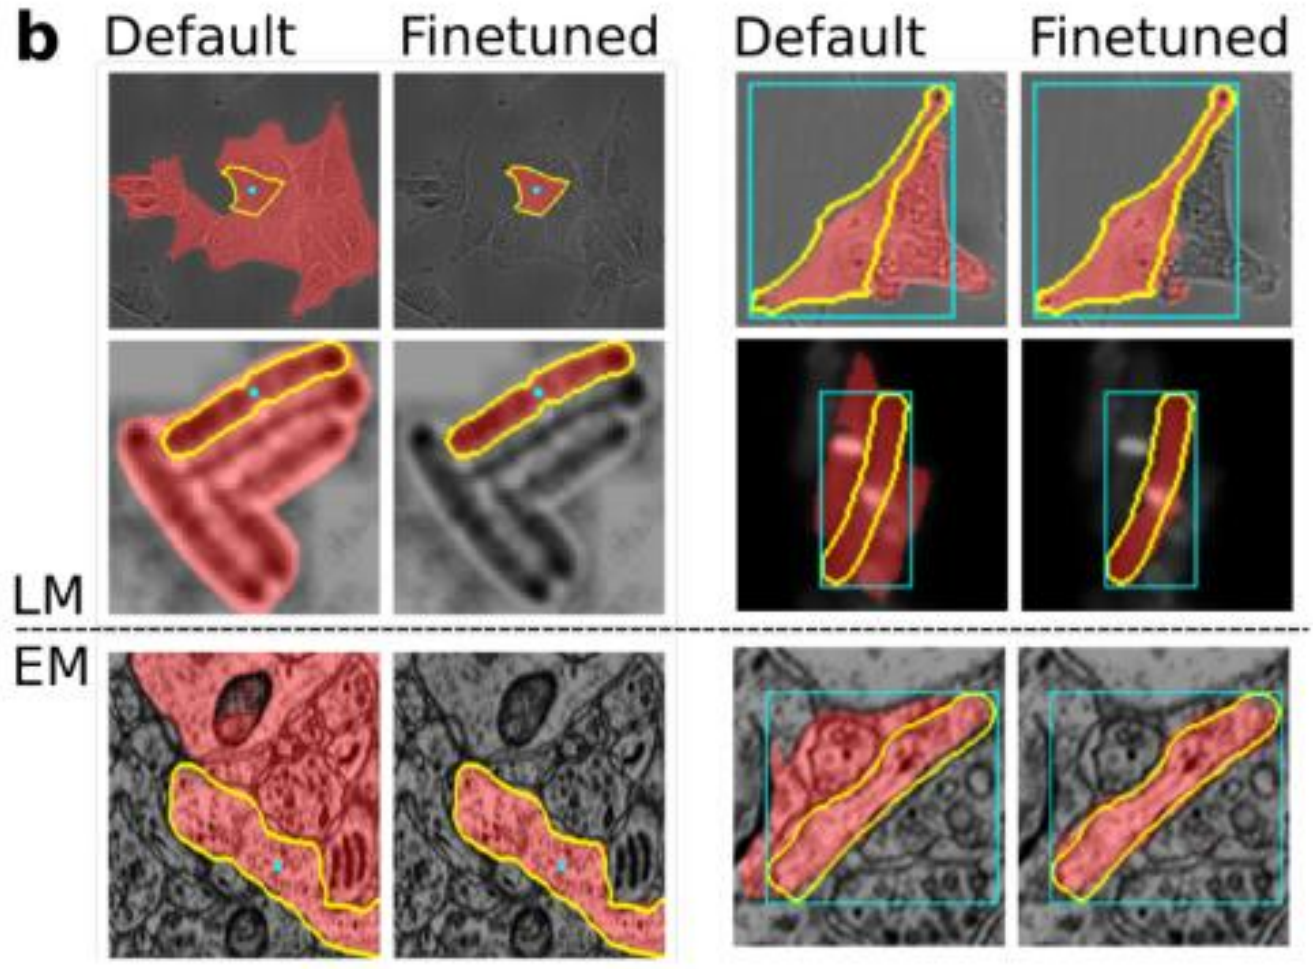

Source: Archit et al (2023) licensed CC-BY 4.0
https://www.biorxiv.org/content/10.1101/2023.08.21.554208v1.article-info

Summary: Deep Learning for Bio-image Analysis

- [Convolutional] Neural Networks is a decade old technology that enabled breakthroughs recently.
  - Image Denoising
  - Image Segmentation
- Common scheme: Smart algorithms for processing input/output of neural networks + standard NN architectures
  - ~~Image in, instance segmentation out~~
- Training these models is
  - computationally expensive,
  - needs large amounts of training data (~~single images~~),
  - requires a certain level of [python] expertise
- If Voronoi-Otsu-Labeling does the job, don't dive into deep learning!

ScaDS.AI
DRESDEN LEIPZIG

CENTER FOR SCALABLE DATA ANALYTICS AND ARTIFICIAL INTELLIGENCE

Exercises

Robert Haase

Funded by

Bundesministerium für Bildung und Forschung

SACHSEN

Diese Maßnahme wird gefördert durch die Bundesregierung aufgrund eines Beschlusses des Deutschen Bundestages. Diese Maßnahme wird mitfinanziert durch Steuermittel auf der Grundlage des von den Abgeordneten des Sächsischen Landtags beschlossenen Haushaltes.

Exercises

Make noise2void, stardist, cellpose [and micro-sam] work.
Hint: This may screw up your conda environment.

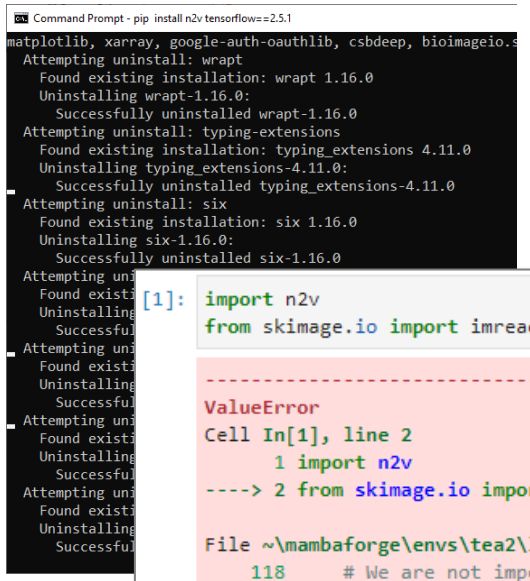

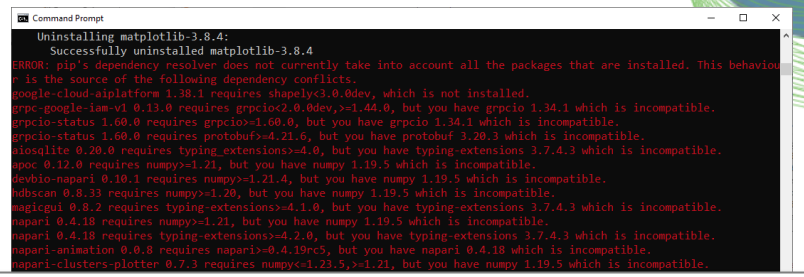

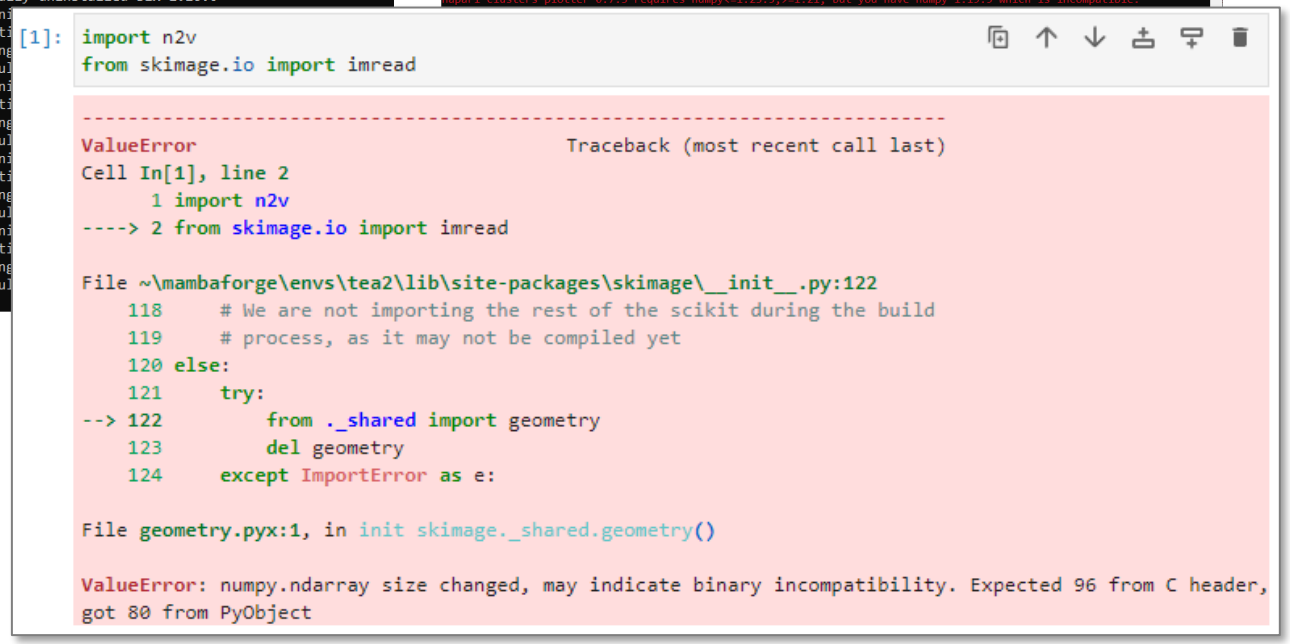

In case of weird errors:

1. Don't panic
2. Recreate you environment
3. Install the thing you want to use, preferably using mamba/conda.

Creating conda environments is not a big deal, just like grabbing a spoon before eating soup.