GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN IN PUBLICA COMMODA
SEIT 1737

COMPUTATIONAL
CELL ANALYTICS

Deep Learning for Image Analysis

Deep Learning Tools for Microscopy

Constantin Pape

Institut für Informatik, Georg August Universität Göttingen

@cppape.bsky.social

https://user.informatik.uni-goettingen.de/~pape41/

Plan for today

- 9:00 - 10:00 - Lecture: Deep Learning Tools, DL State-of-the-art, Data Annotation
- 10:00 - ca. 12:00 - Exercise 3
- ca. 12:00 - 13:00 - Lunch
- ca. 13:00 - 14:30 - Exercise 3
- 14:30 - 15:30 - Project presentations
- 15:30 - 17:00 - Project planning and getting started

Tools for Segmentation & Other Tasks

ilastik

https://www.ilastik.org/

Software for bioimage segmentation based on **“classical” machine learning**: pre-defined features + shallow classifier

Different workflows:

- Pixel Classification (= semantic segmentation)
- Object Classification (= (cell) classification)
- Carving (interactive 3D instance segmentation)
- …
- Prediction with pre-trained networks from BioImage.IO (no training)

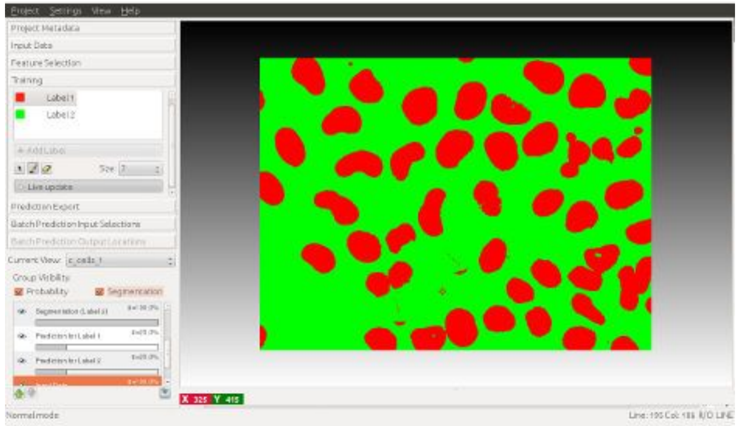

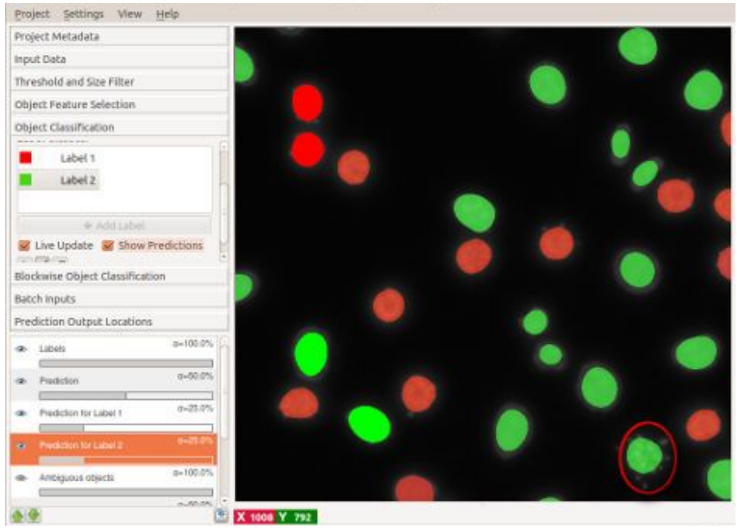

ilastik

https://www.ilastik.org/

Software for bioimage segmentation based on **“classical” machine learning**: pre-defined features + shallow classifier

Different workflows:

- Pixel Classification (= semantic segmentation)
- Object Classification (= (cell) classification)

Stand-alone application, written in python

Strengths:

- Interactive training with few “brushstroke” labels
- Easy to install and use, does not need a GPU

Limitations:

- Not powerful enough to solve challenging tasks, e.g. cell instance segmentation with complex shapes

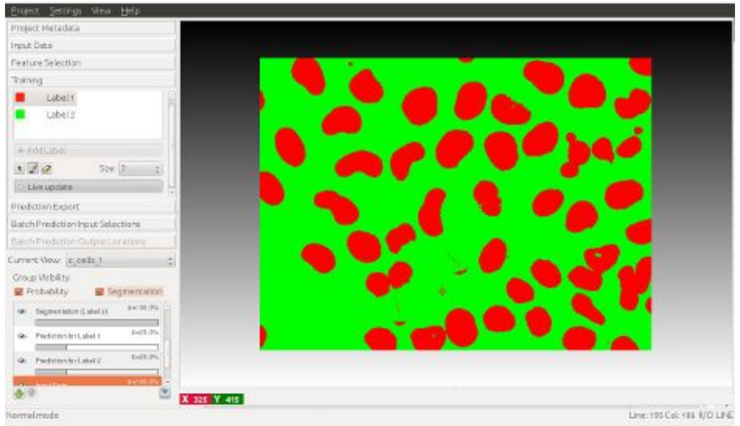

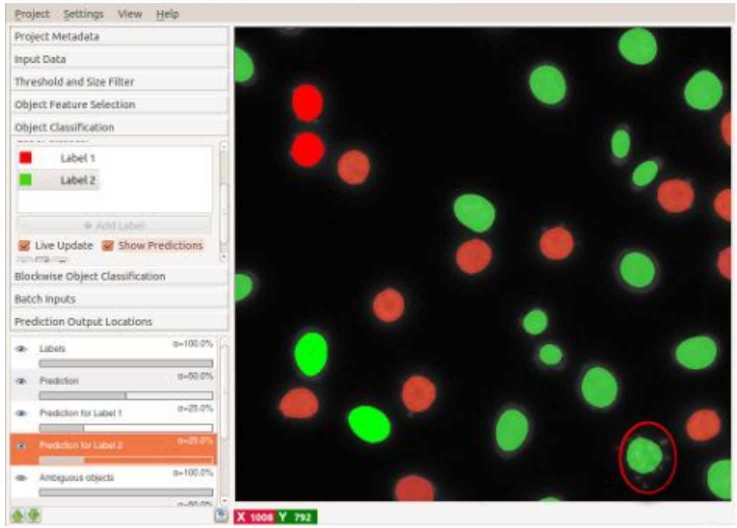

Deep learning-based tools: “(Almost) everything is a U-Net”

U-Net is the working horse of bioimage analysis. Most deep learning tools implement (and pretrain) a U-Net version for a specific task.

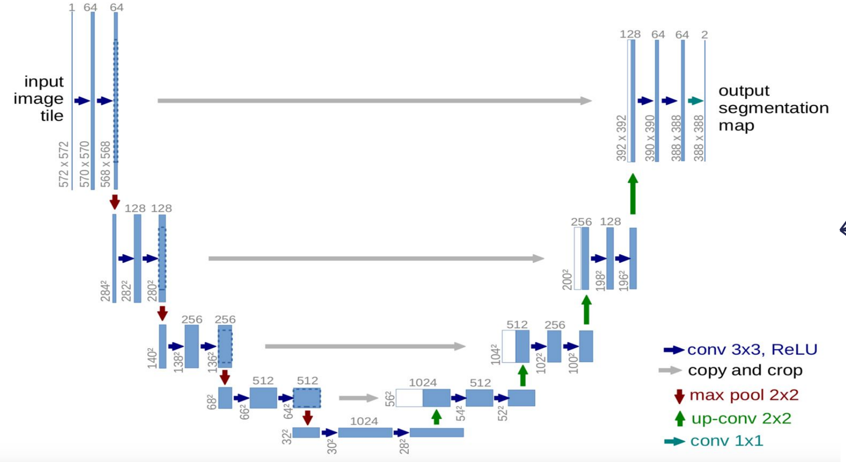

Cell & Nucleus Segmentation

StarDist

Cellpose

CochleaNet

EM Segmentation

MitoNet

MemBrain

SynapseNet

Spot Detection

Spotiflow

STACC

Denoising / Image Restoration

CAREamics

General-purpose Segmentation: nn-UNet

U-Net implementation that selects good architecture settings based on the data characteristics and automatically optimizes other settings in training

Very popular in medical image analysis.

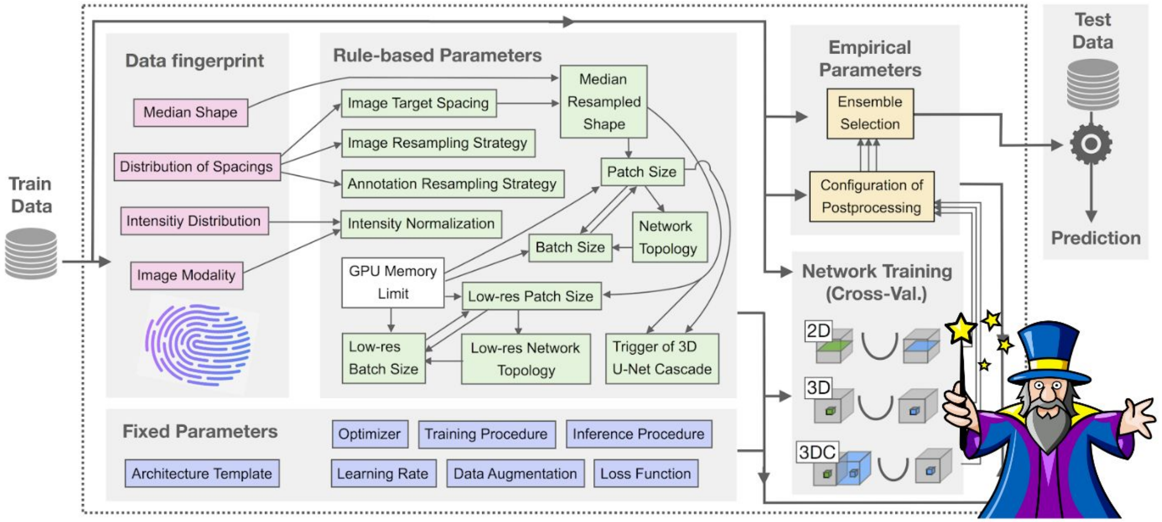

Train Data

Data fingerprint

Median Shape

Distribution of Spacings

Intensitiy Distribution

Image Modality

Rule-based Parameters

Image Target Spacing

Image Resampling Strategy

Annotation Resampling Strategy

Intensity Normalization

GPU Memory Limit

Low-res Batch Size

Median Resampled Shape

Patch Size

Batch Size

Network Topology

Low-res Patch Size

Low-res Network Topology

Trigger of 3D U-Net Cascade

Fixed Parameters

Architecture Template

Optimizer

Training Procedure

Inference Procedure

Learning Rate

Data Augmentation

Loss Function

Empirical Parameters

Ensemble Selection

Configuration of Postprocessing

Network Training (Cross-Val.)

2D

3D

3DC

Test Data

Prediction

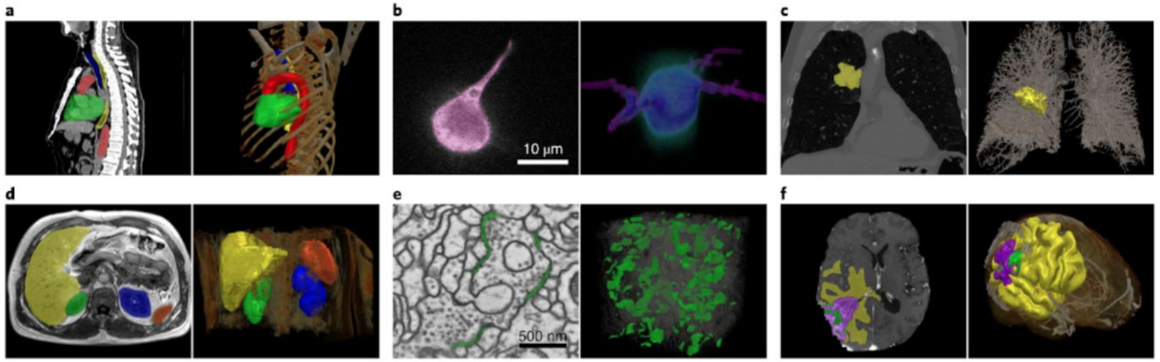

https://github.com/MIC-DKFZ/nnUNet

General-purpose Segmentation: nn-UNet

U-Net implementation that selects good architecture settings based on the data characteristics and automatically optimizes other settings in training

Very popular in medical image analysis.

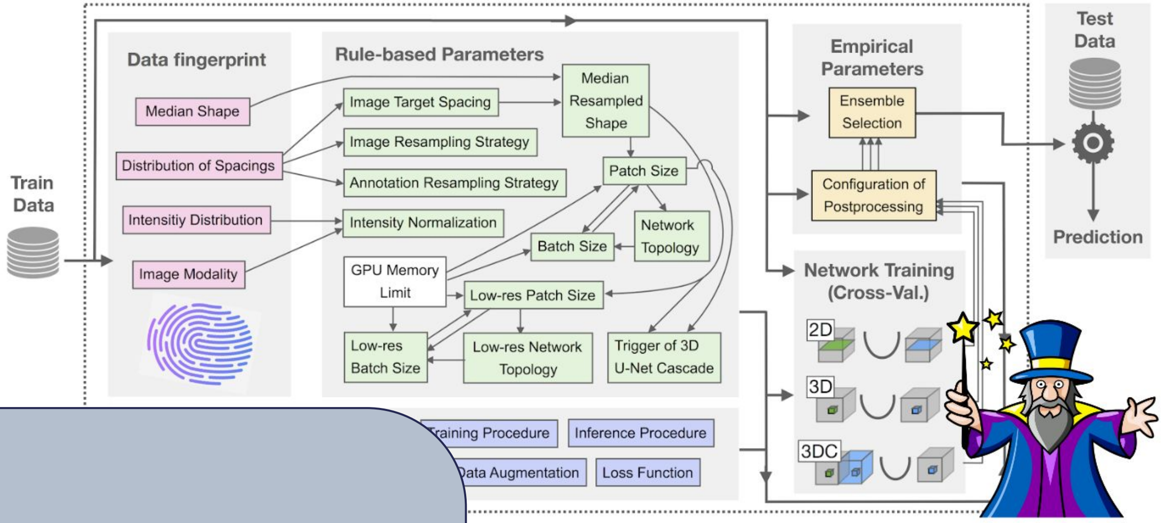

Python library, based on PyTorch

Strengths:

- Training and prediction based on configuration files
- High-quality results due to many optimizations

Limitations:

- No pretrained models -> always need labels
- Only semantic segmentation, no (direct) instance segmentation

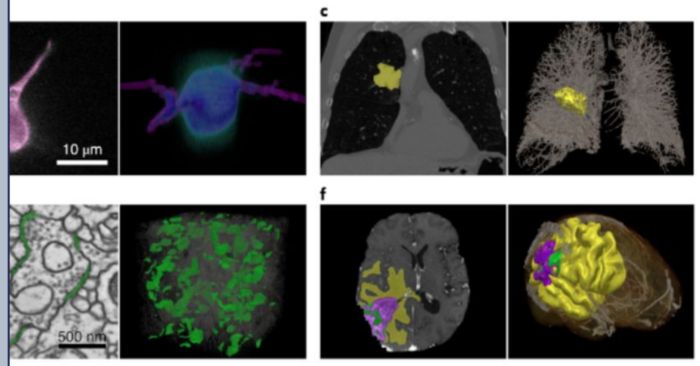

https://github.com/MIC-DKFZ/nnUNet

Nucleus Segmentation: StarDist

https://github.com/stardist/stardist

Tool for instance segmentation of (star-)convex objects in 2D and 3D. **Popular for nucleus segmentation.**

Based on U-Net that predicts center distances across rays.

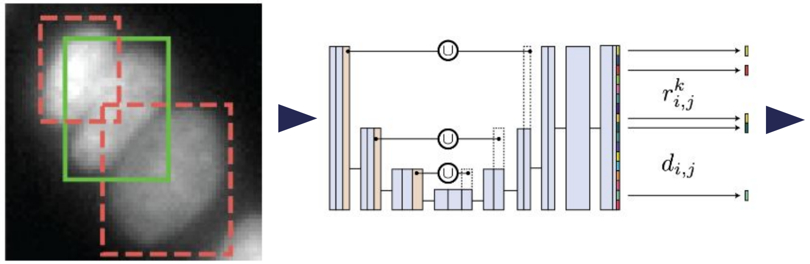

Prediction:

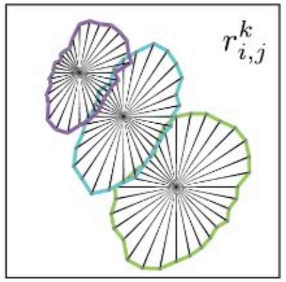

Center Distances

\+ Foreground

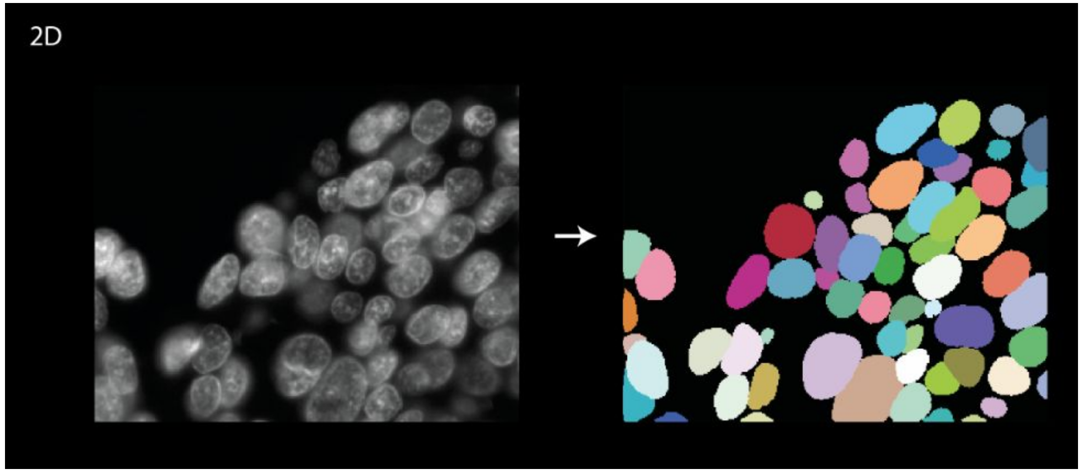

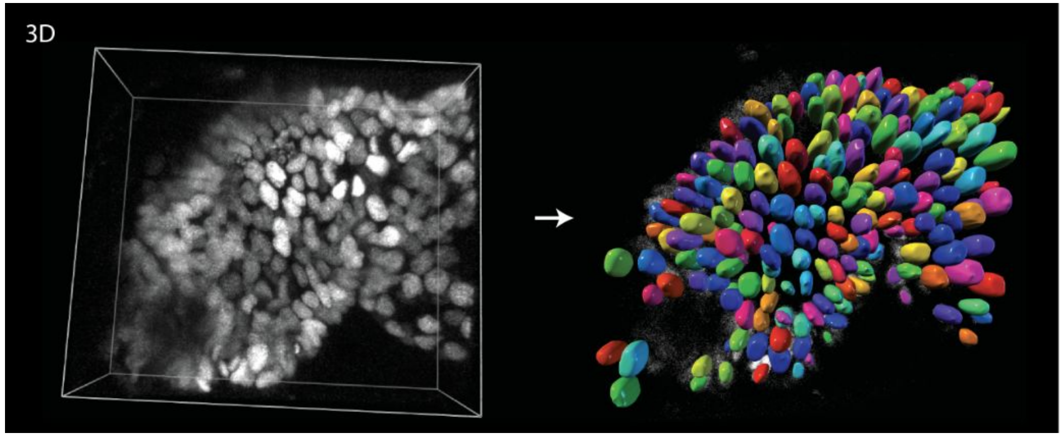

Nucleus Segmentation: StarDist

https://github.com/stardist/stardist

Tool for instance segmentation of (star-)convex objects in 2D and 3D. **Popular for nucleus segmentation.**

Based on U-Net that predicts center distances across rays.

Prediction:

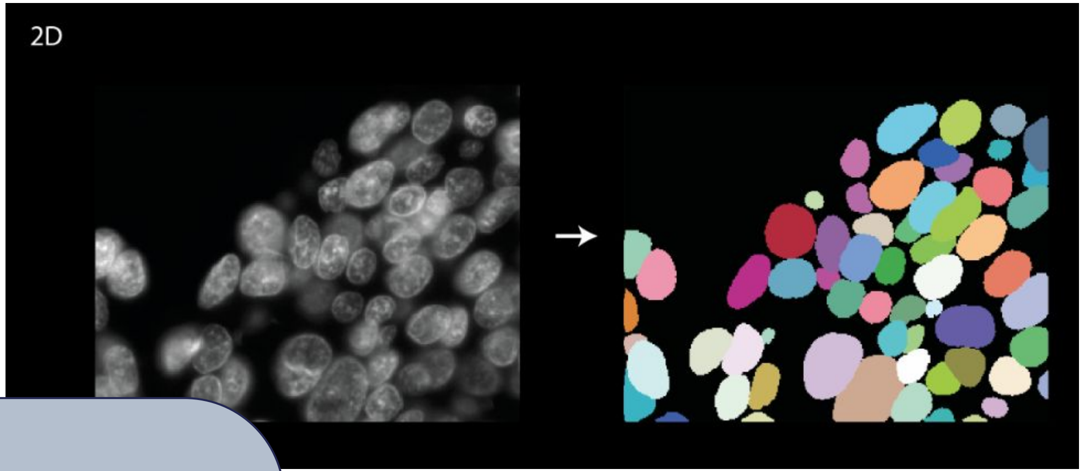

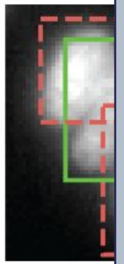

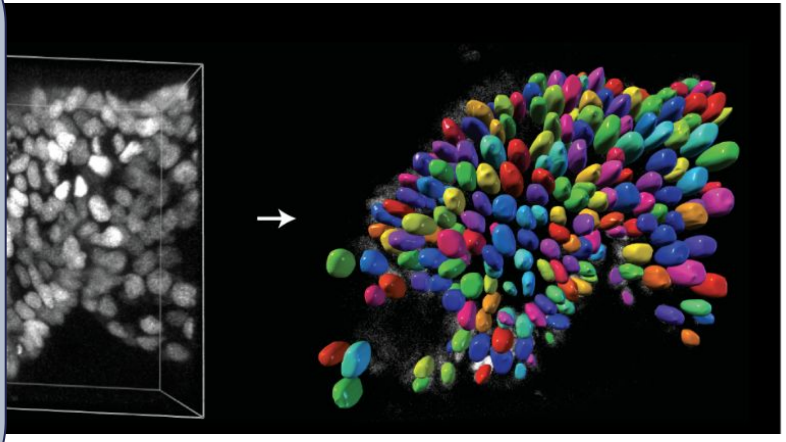

Python library, based on Tensorflow & napari-plugin

Strengths:

- Very good pre-trained models for nucleus segmentation in fluorescence microscopy + H&E staining

Limitations:

- No good pretrained model for 3D available
- Does not work well for complex shapes (not star-convex)

Cell & Nucleus Segmentation: CellPose

Most popular method for cell segmentation

Based on U-Net that predicts gradient flow to cell center

- Pre-trained models for cell and nucleus segmentation in light microscopy
- Combined model for segmentation & denoising (Cellpose 3)
- Human-in-the loop training for fast data annotation (details later)
- CellposeSAM based on vision transformer, similar to µSAM (details later)

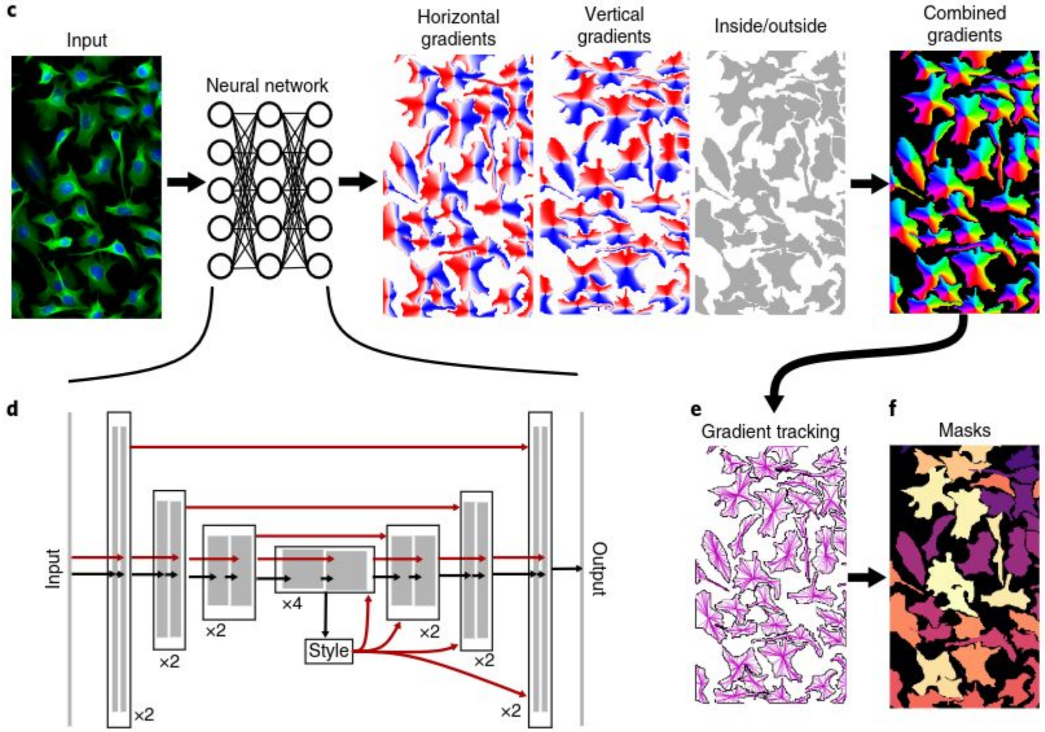

https://www.cellpose.org/

Cell & Nucleus Segmentation: CellPose

Most popular method for cell segmentation

Based on U-Net that predicts gradient flow to cell center

- Pre-trained models for cell and nucleus segmentation in light microscopy
- Combined model for segmentation &

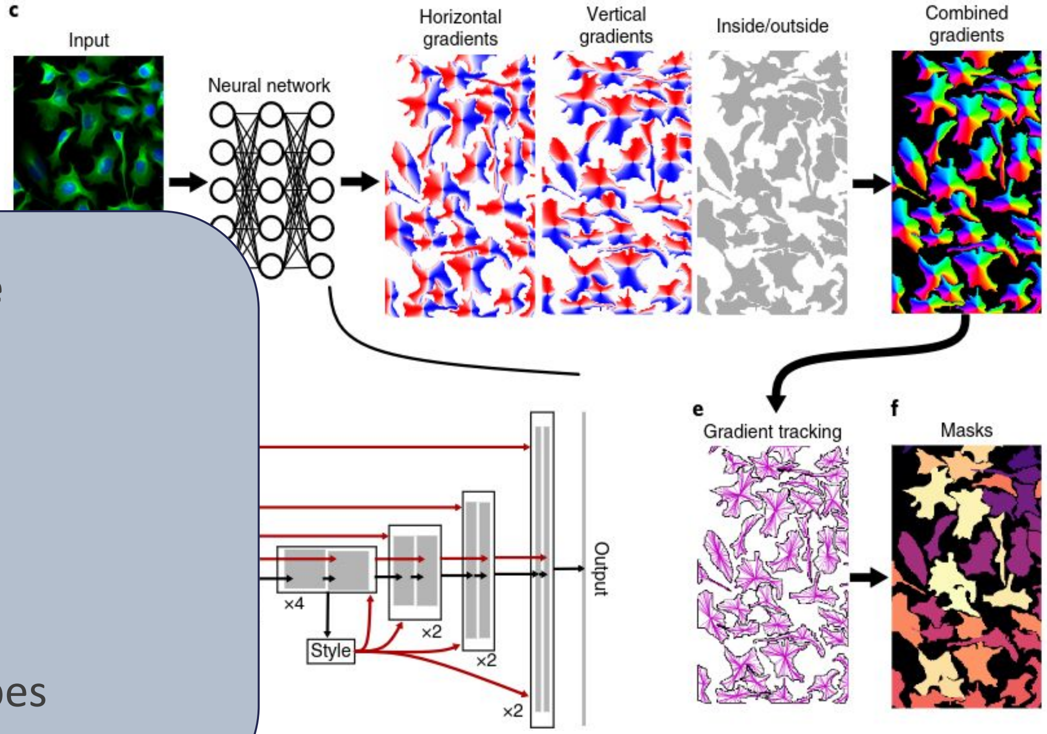

Python library, based on PyTorch & stand-alone software

Strengths:

- Very good pre-trained models for cell segmentation
- Fast data annotation via human-in-the loop training

Limitations:

- No real 3D support, sub-optimal for complex 3D shapes

https://www.cellpose.org/

Mitochondrion Segmentation (EM): MitoNet

Model for mitochondrion instance segmentation

in (volume) EM

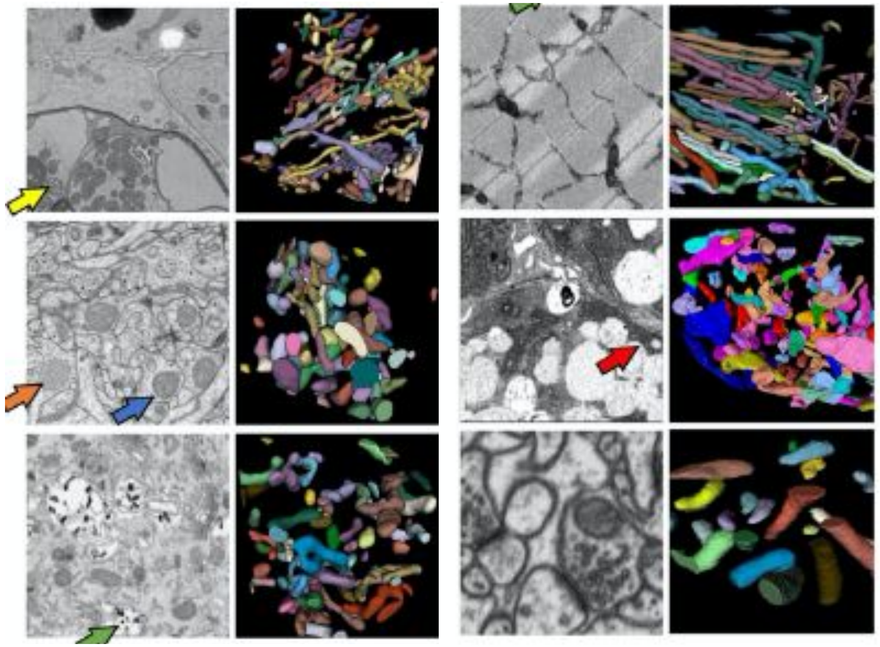

Python library, based on PyTorch & napari plugin

Strengths:

- Good pre-trained model for mito segmentation
- Supports 2D and 3D data

https://volume-em.github.io/empanada.html

Membrane segmentation in Cryo-ET: MemBrain

Pertained network (based on nn U-Net) for membrane segmentation in CryoET

- Particle picking on membranes

Python library, based on PyTorch & napari plugin

Strengths:

- Good pre-trained model for membrane segmentation

Limitations:

- Only for CryoET data

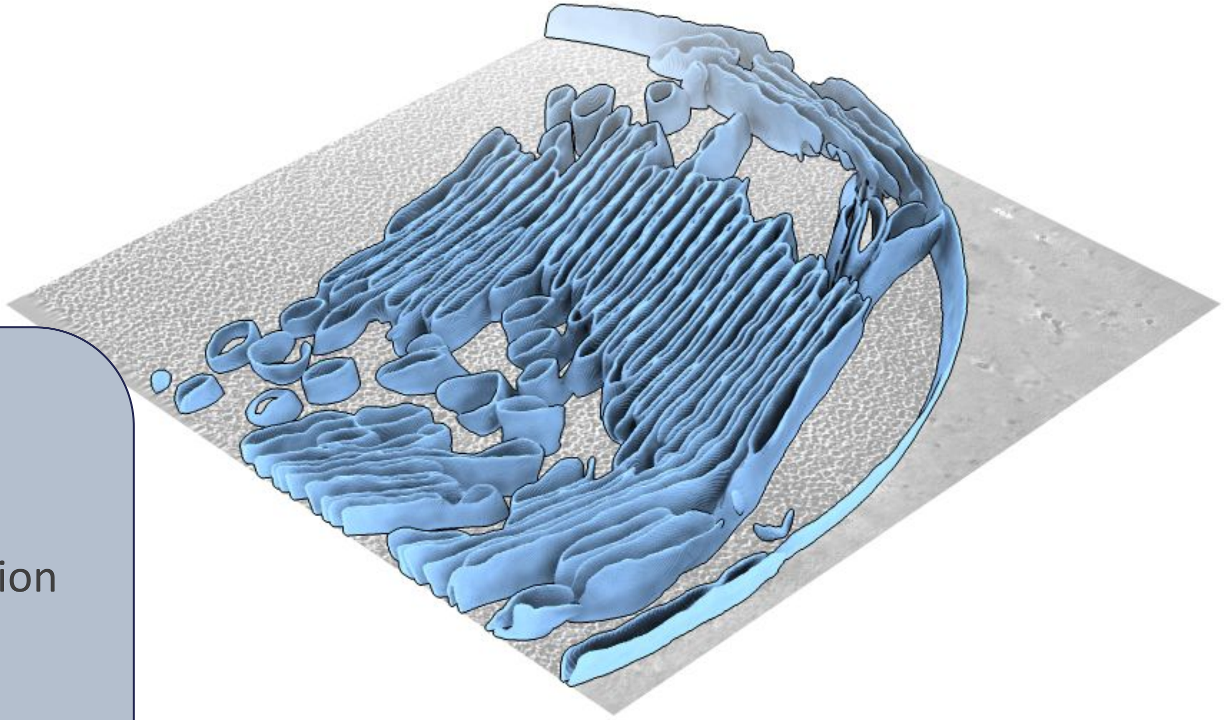

https://github.com/CellArchLab/MemBrain-v2

Segmentation of Synaptic Structures (EM): SynapseNet

Tool from my group, Cooper group & others for segmentation of vesicles and other structures in electron microscopy, esp. electron tomography

Domain adaptation enables transfer to other data without additional labels

Python library, based on PyTorch & napari plugin

Strengths:

- Good pre-trained model for vesicle segmentation in 2D & 3D EM
- Networks for mitos, synaptic compartments and active zone (in 3D)

Limitations:

- Domain adaptation only works if initial network predictions are of sufficient quality

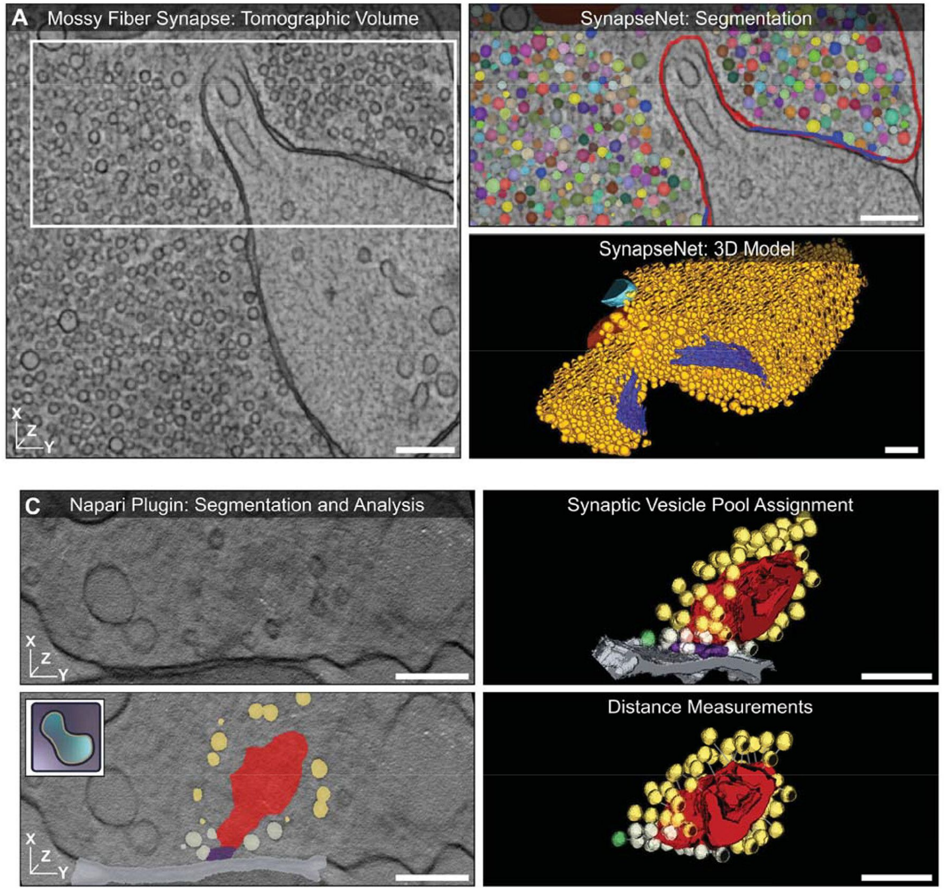

https://github.com/computational-cell-analytics/synapse-net

Segmentation of large light sheet data (of the cochlea): CochleaNet

Tool from my group & Moser group to segment cells in light-sheet microscopy (LSM) of the cochlea.

Extension to general-purpose cell segmentation in LSM planned.

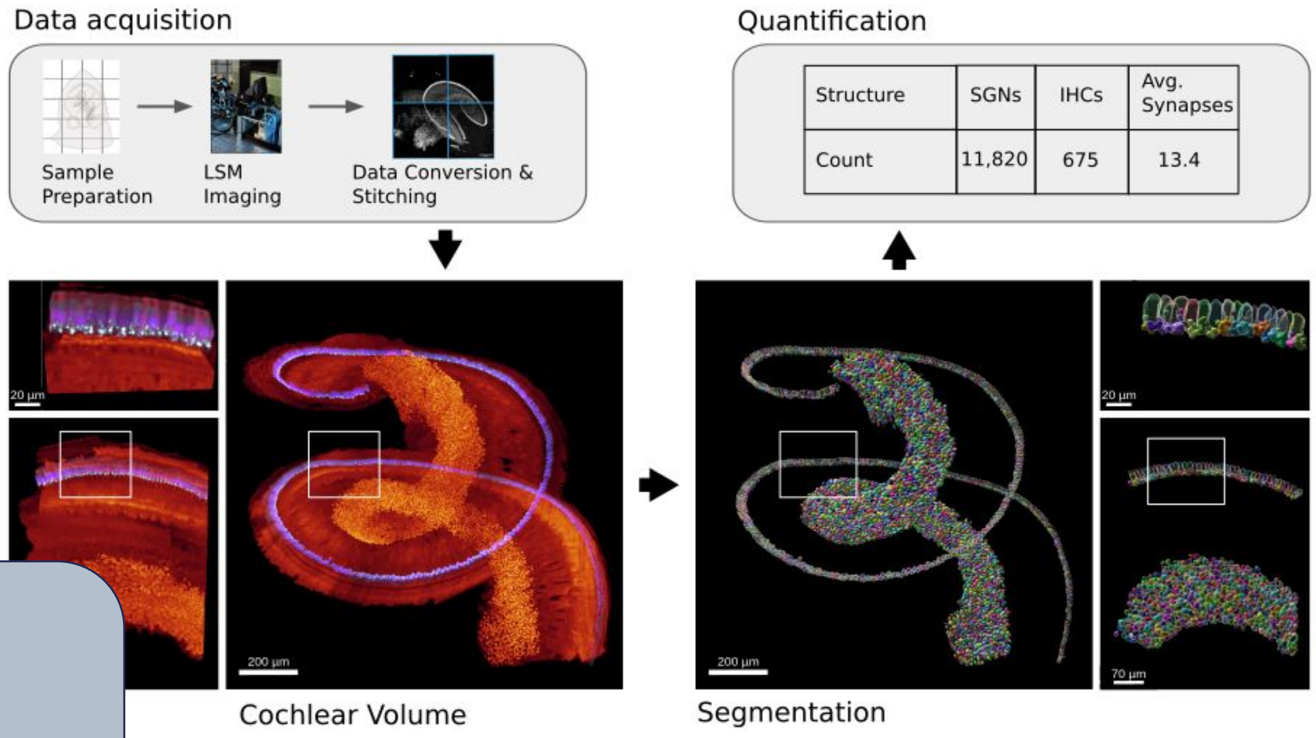

Structure	SGNs	IHCs	Avg. Synapses
Count	11,820	675	13.4

Python library, under development!
Strengths:

- Can segment very large volumetric data.

Limitations:

- Under development.
- Currently specific to cochlear cell types.

https://github.com/computational-cell-analytics/cochlea-net

Spot Detection

Spotiflow

Detect coordinates of spots (small bright structures) in images

SpotiFlow: General purpose model for spot detection in fluorescence microscopy

STACC: Tool from my group for bacterial colony counting based using a similar approach

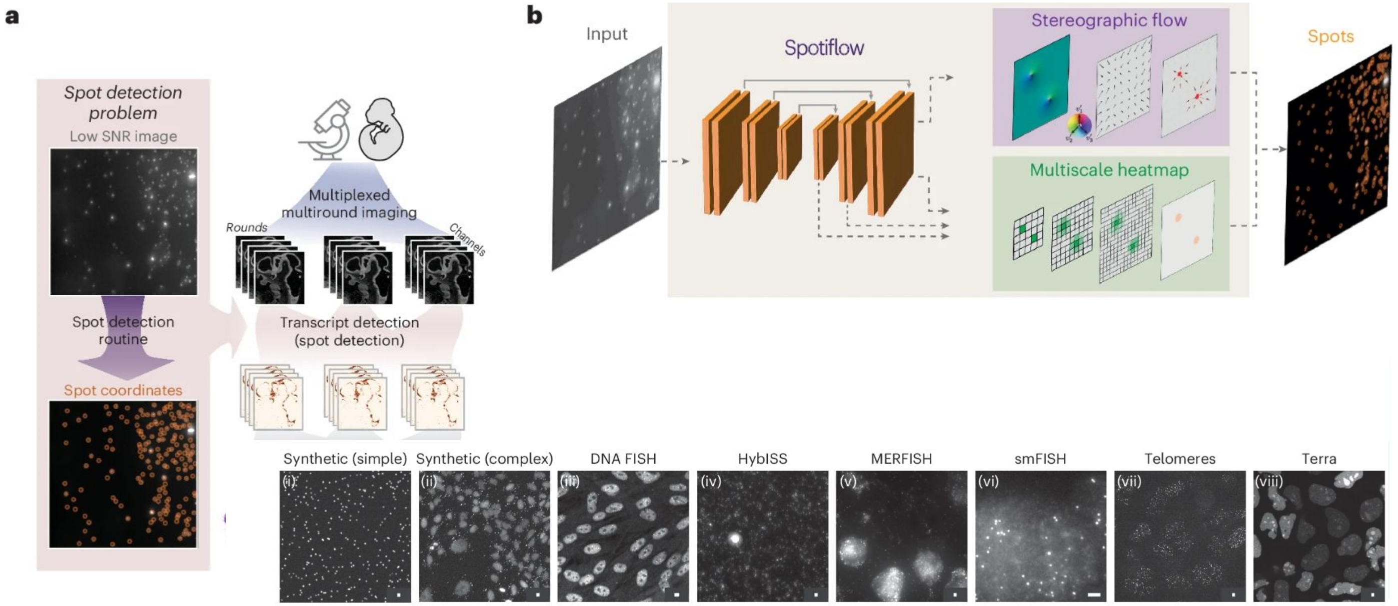

https://github.com/weigertlab/spotiflow

https://github.com/computational-cell-analytics/stacc

Cell Tracking

“Classical tools”:

- TrackMate: https://imagej.net/plugins/trackmate/
- btrack: https://github.com/quantumjot/btrack
- motile: https://github.com/funkelab/motile

Deep-learning based:

- Trackastra: https://github.com/weigertlab/trackastra

See lecture yesterday for details.

Denoising / Image Restoration

Library for different denoising and image restoration tasks: CAREamics

https://careamics.github.io

See lecture yesterday for details.

Why we need a shared “modelzoo”!

https://bioimage.io

A lot of different models for different applications.

Goal of BioImage.IO (recall lecture yesterday):

- A central place where you can search for models for your problem.
- Standardized format to support multiple tools.

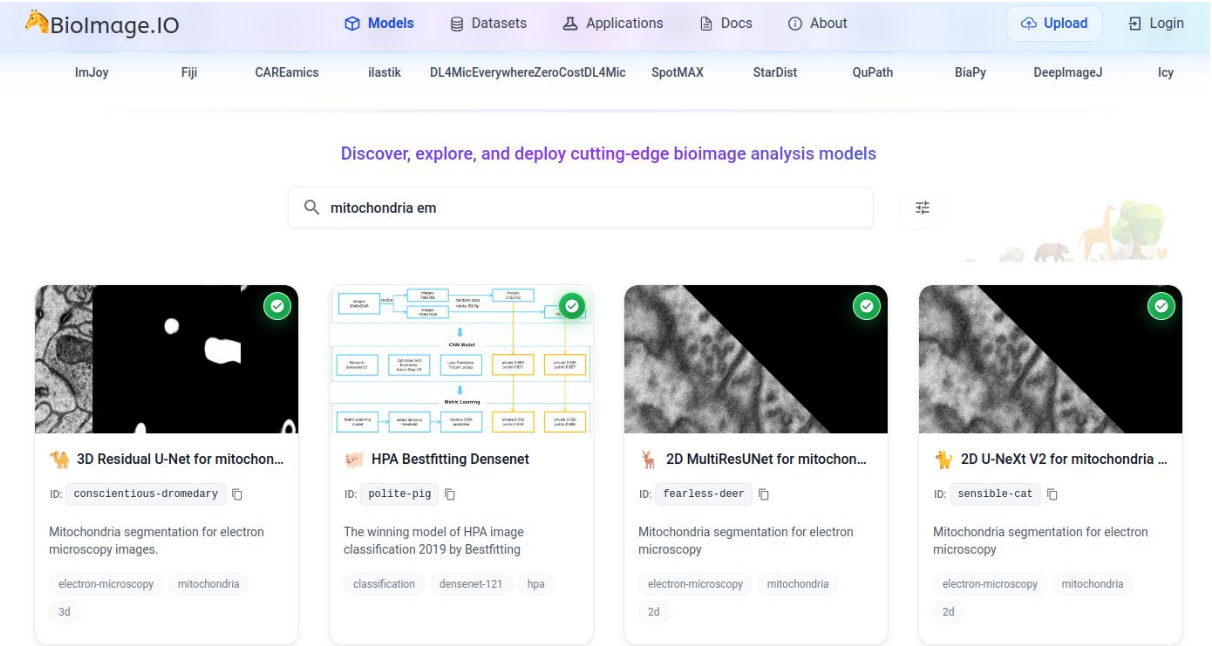

BioImage.IO — Models — Datasets — Applications — Docs — About — Upload — Login

ImJoy — Fiji — CAREamics — ilastik — DL4MicEverywhereZeroCostDL4Mic — SpotMAX — StarDist — QuPath — BiaPy — DeepImageJ — Icy

Discover, explore, and deploy cutting-edge bioimage analysis models

mitochondria em

3D Residual U-Net for mitochon...
ID: conscientious-dromedary
Mitochondria segmentation for electron microscopy images.
electron-microscopy · mitochondria · 3d

HPA Bestfitting Densenet
ID: polite-pig
The winning model of HPA image classification 2019 by Bestfitting
classification · densenet-121 · hpa

2D MultiResUNet for mitochon...
ID: fearless-deer
Mitochondria segmentation for electron microscopy
electron-microscopy · mitochondria · 2d

2D U-NeXt V2 for mitochondria ...
ID: sensible-cat
Mitochondria segmentation for electron microscopy
electron-microscopy · mitochondria · 2d

Tools / Software Supporting BioImage.IO

ilastik

BiaPy

#ZeroCostDL4Mic

deepImageJ

QuPath

Other Libraries / Frameworks / Tools

- ZeroCostDL4MIC: Jupyter Notebooks implementing many different deep learning methods
  - https://github.com/HenriquesLab/ZeroCostDL4Mic

- BiaPy: PyTorch-based library for deep learning in microscopy
  - https://github.com/BiaPyX/BiaPy

- torch-em: My library for deep learning in microscopy
  - https://github.com/constantinpape/torch-em

- PyTorch Lightning: Library that simplifies the use of PyTorch (general-purpose)
  - https://lightning.ai/docs/pytorch/stable/

State-of-the-art DL / AI

Large Language Models

Large Language Models

Architecture underlying ChatGPT & Co:

- *Transformer* architecture: many *self-attention* layers
- model for processing sequences with interactions between all elements in the sequence

Large Language Models

Architecture underlying ChatGPT & Co:

- model for processing sequences with interactions between all elements in the sequence
- *Transformer* architecture: many *self-attention* layers

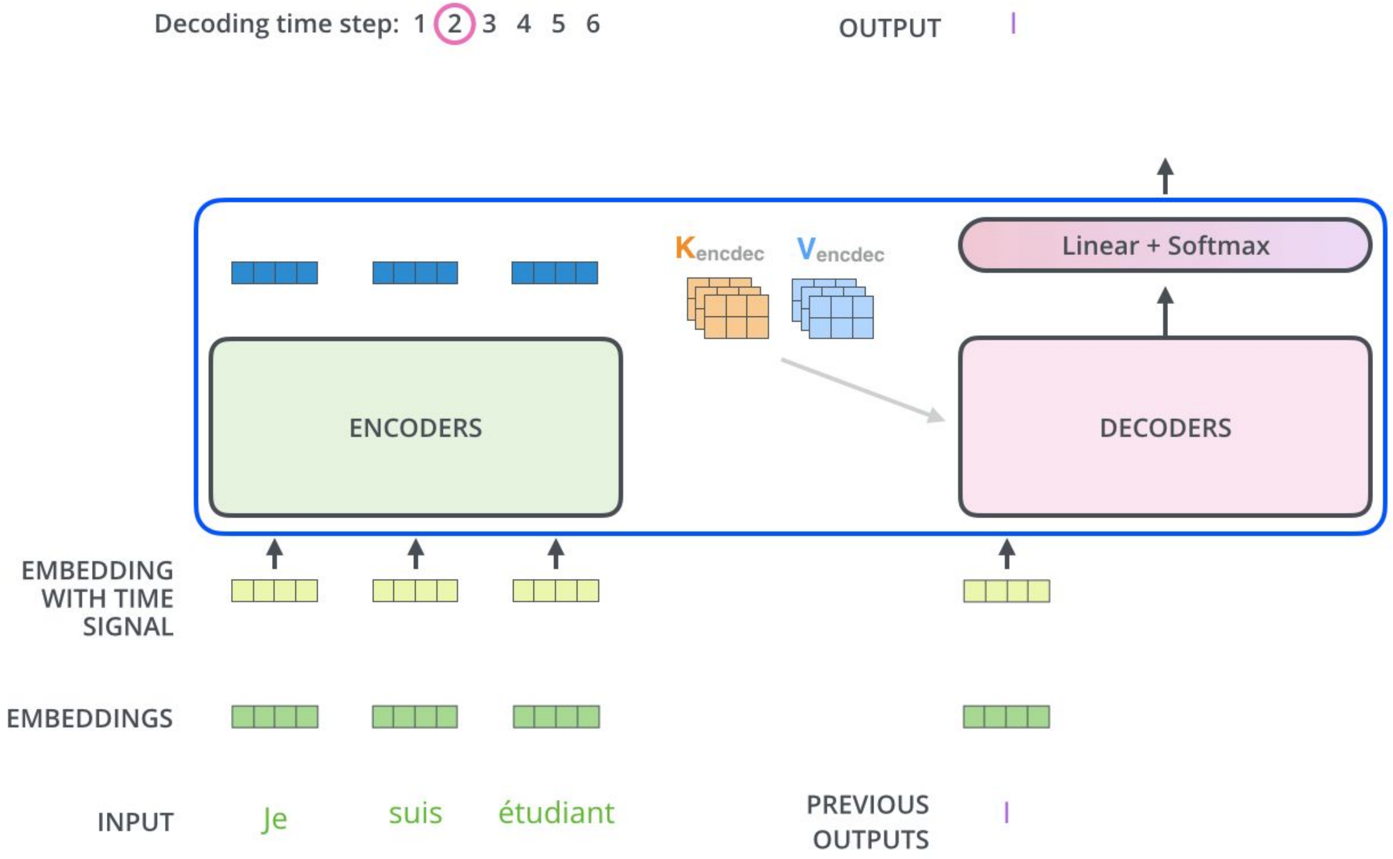

Large Language Models

Architecture underlying ChatGPT & Co:

- model for processing sequences with interactions between all elements in the sequence

- *Transfor* e: many *se* rs

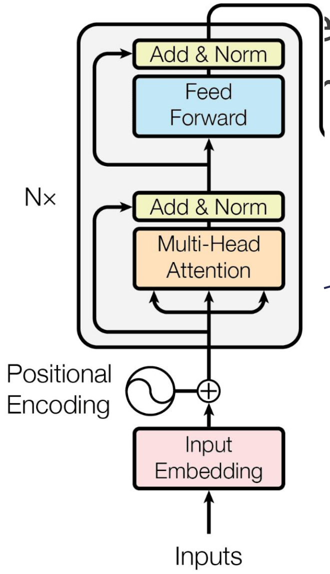

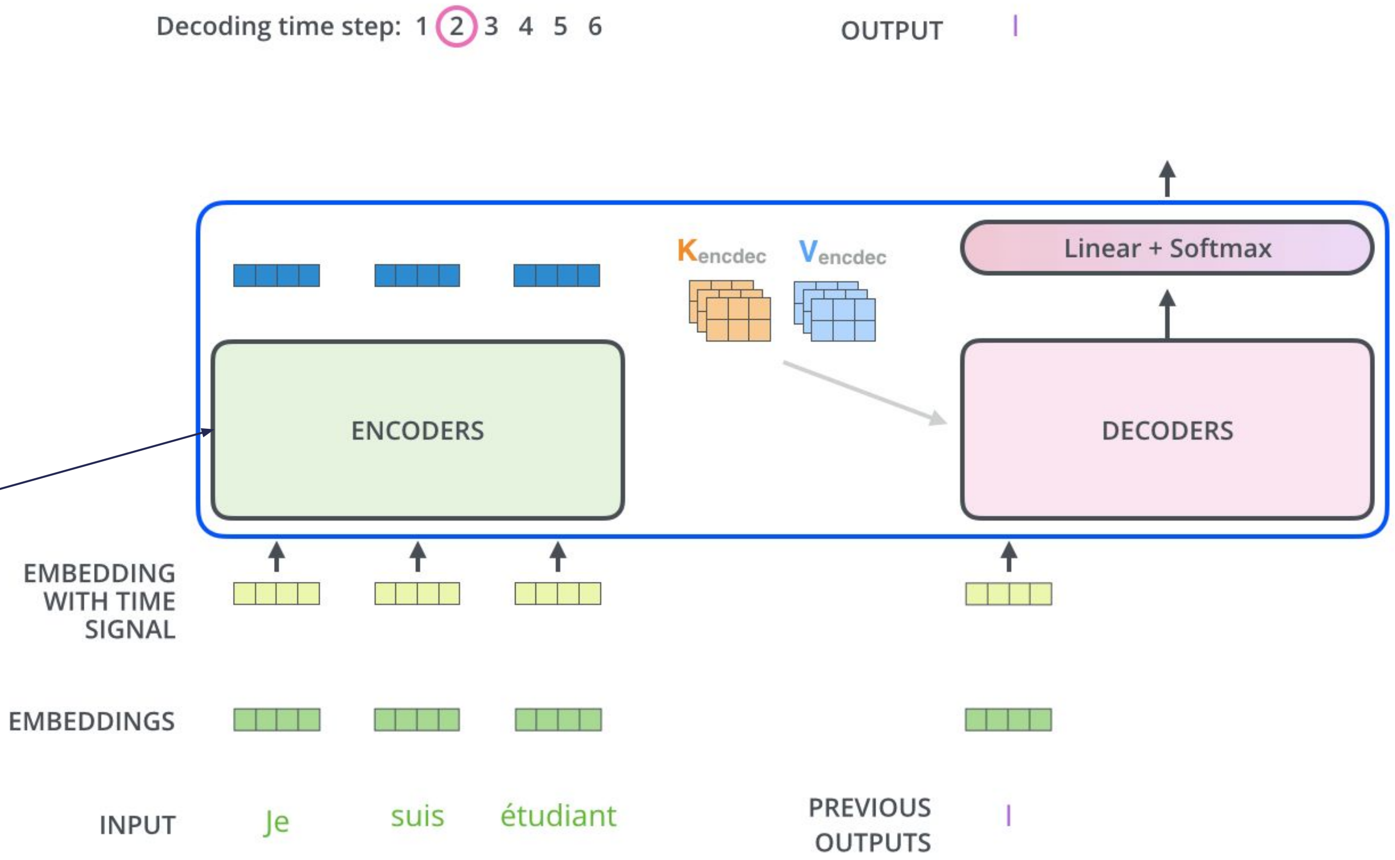

Large Language Models

- Transformers for translation: supervised training with paired text

Large Language Models

- Transformers for translation: supervised training with paired text

- Chatbots / Chat GPT:
  - unsupervised (self-supervised) training: predict next word in sequence

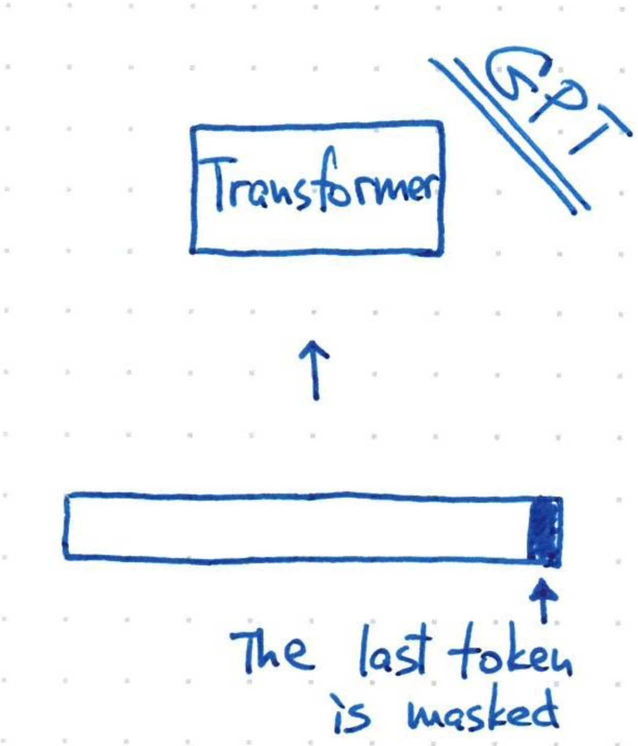

Large Language Models

- Transformers for translation: supervised training with paired text

- Chatbots / Chat GPT:
  - unsupervised (self-supervised) training: predict next word in sequence
  - trained on very large dataset
  - followed by supervised fine tuning / reinforcement learning

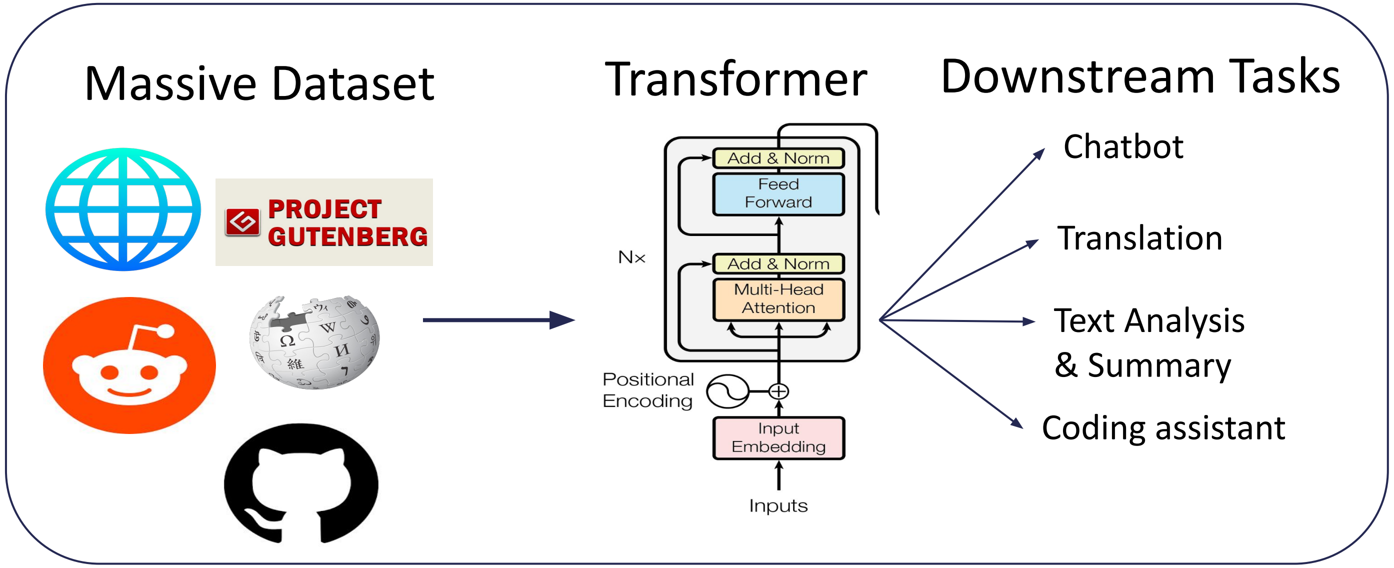

Massive Dataset

Transformer

Downstream Tasks

- Chatbot
- Translation
- Text Analysis & Summary
- Coding assistant

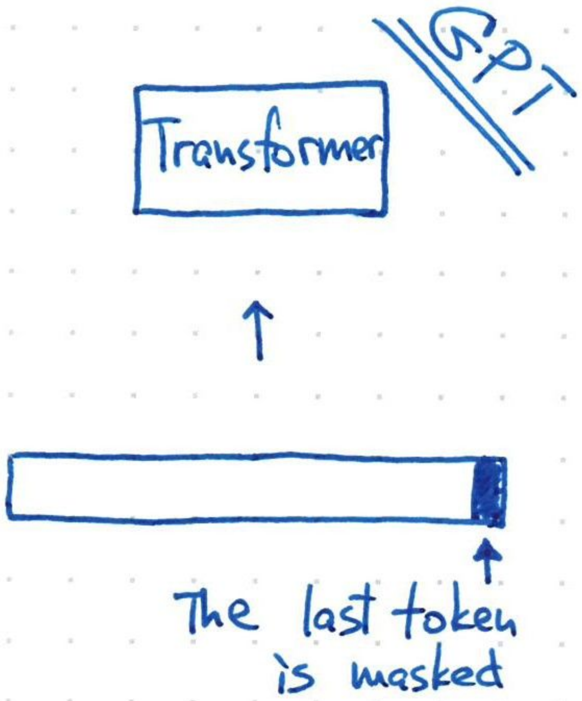

State-of-the-art DL / AI

Vision Transformer or “Not everything is a U-Net anymore”

Vision Transformer

Cut up image into patches

Order patches sequentially

Process with transformer architecture

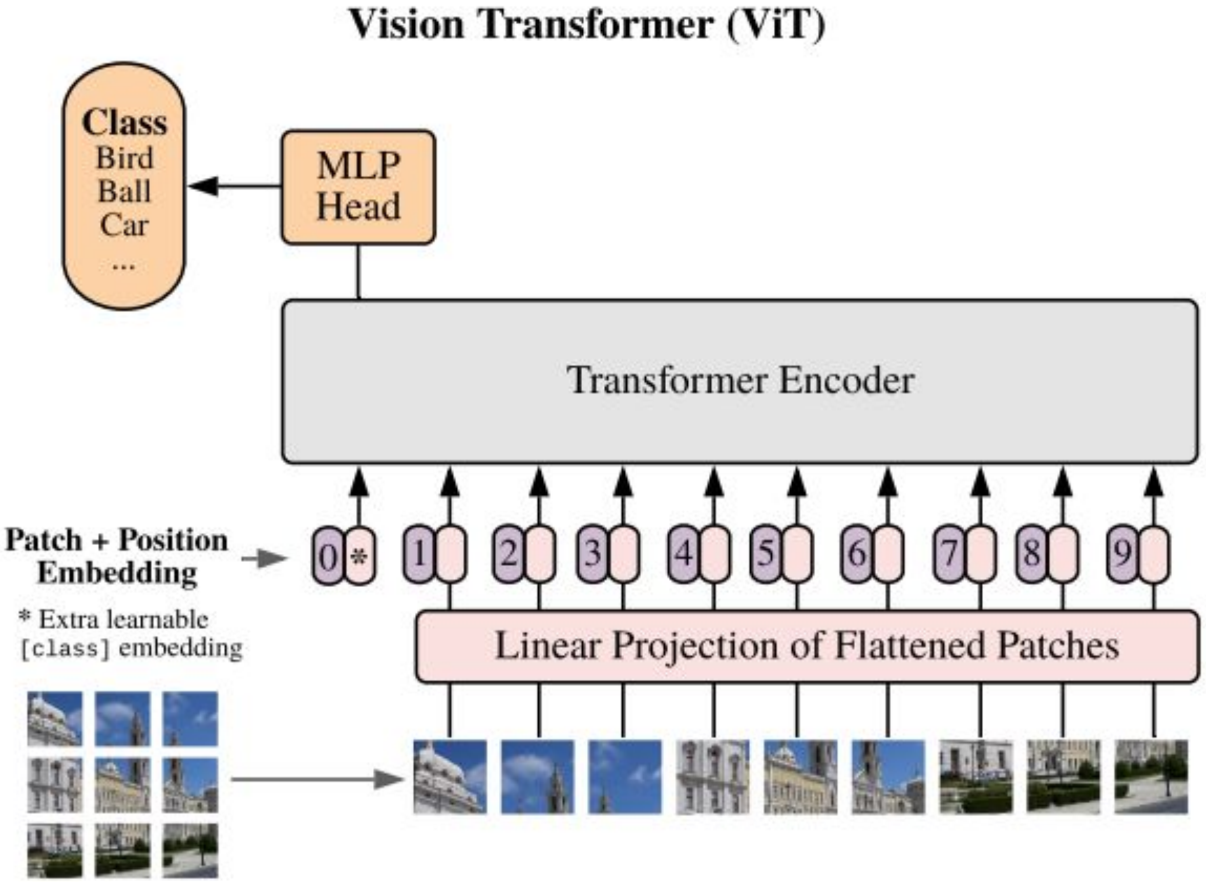

Vision Transformer

Cut up image into patches

Order patches sequentially

Process with transformer architecture

Classification: combine sequence representation at the end and add small multi-layer network for classification

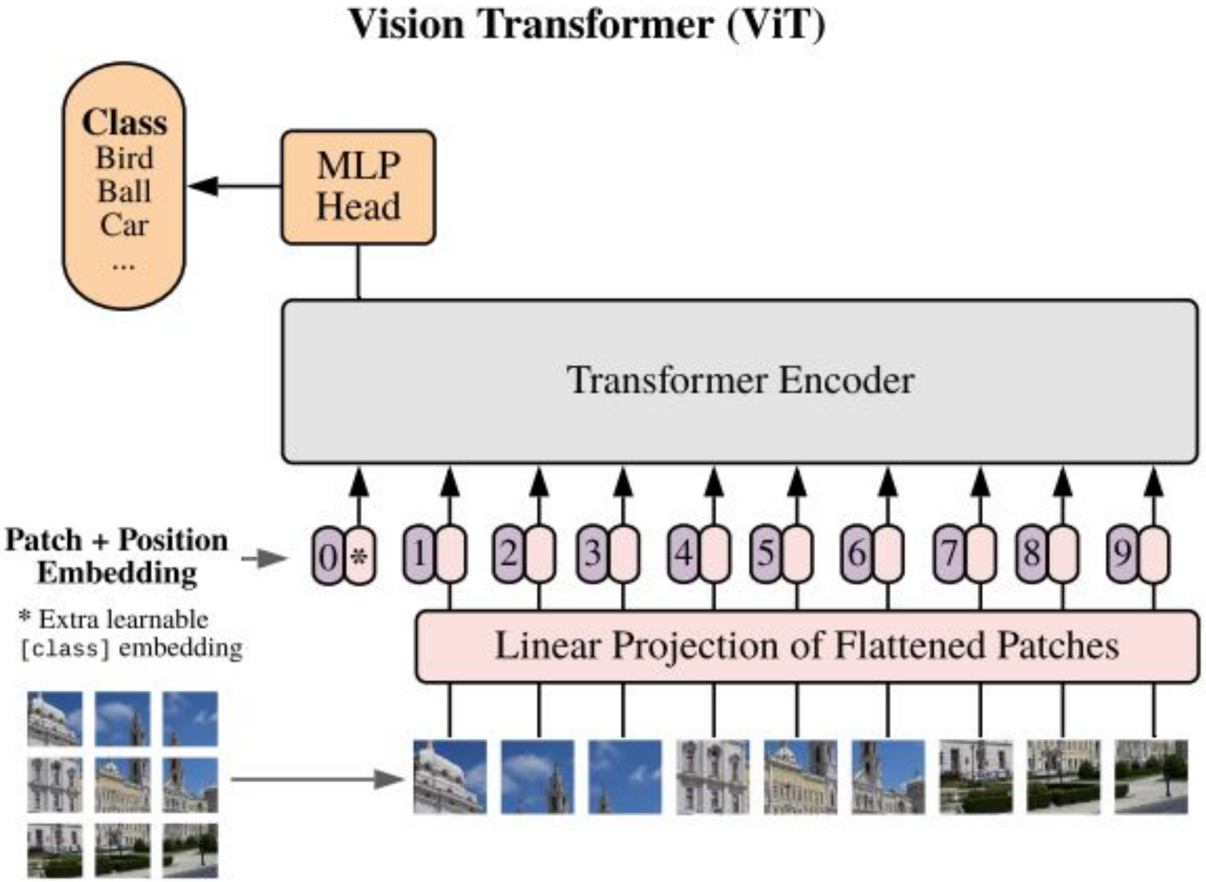

Segmentation: U-Net inspired architecture

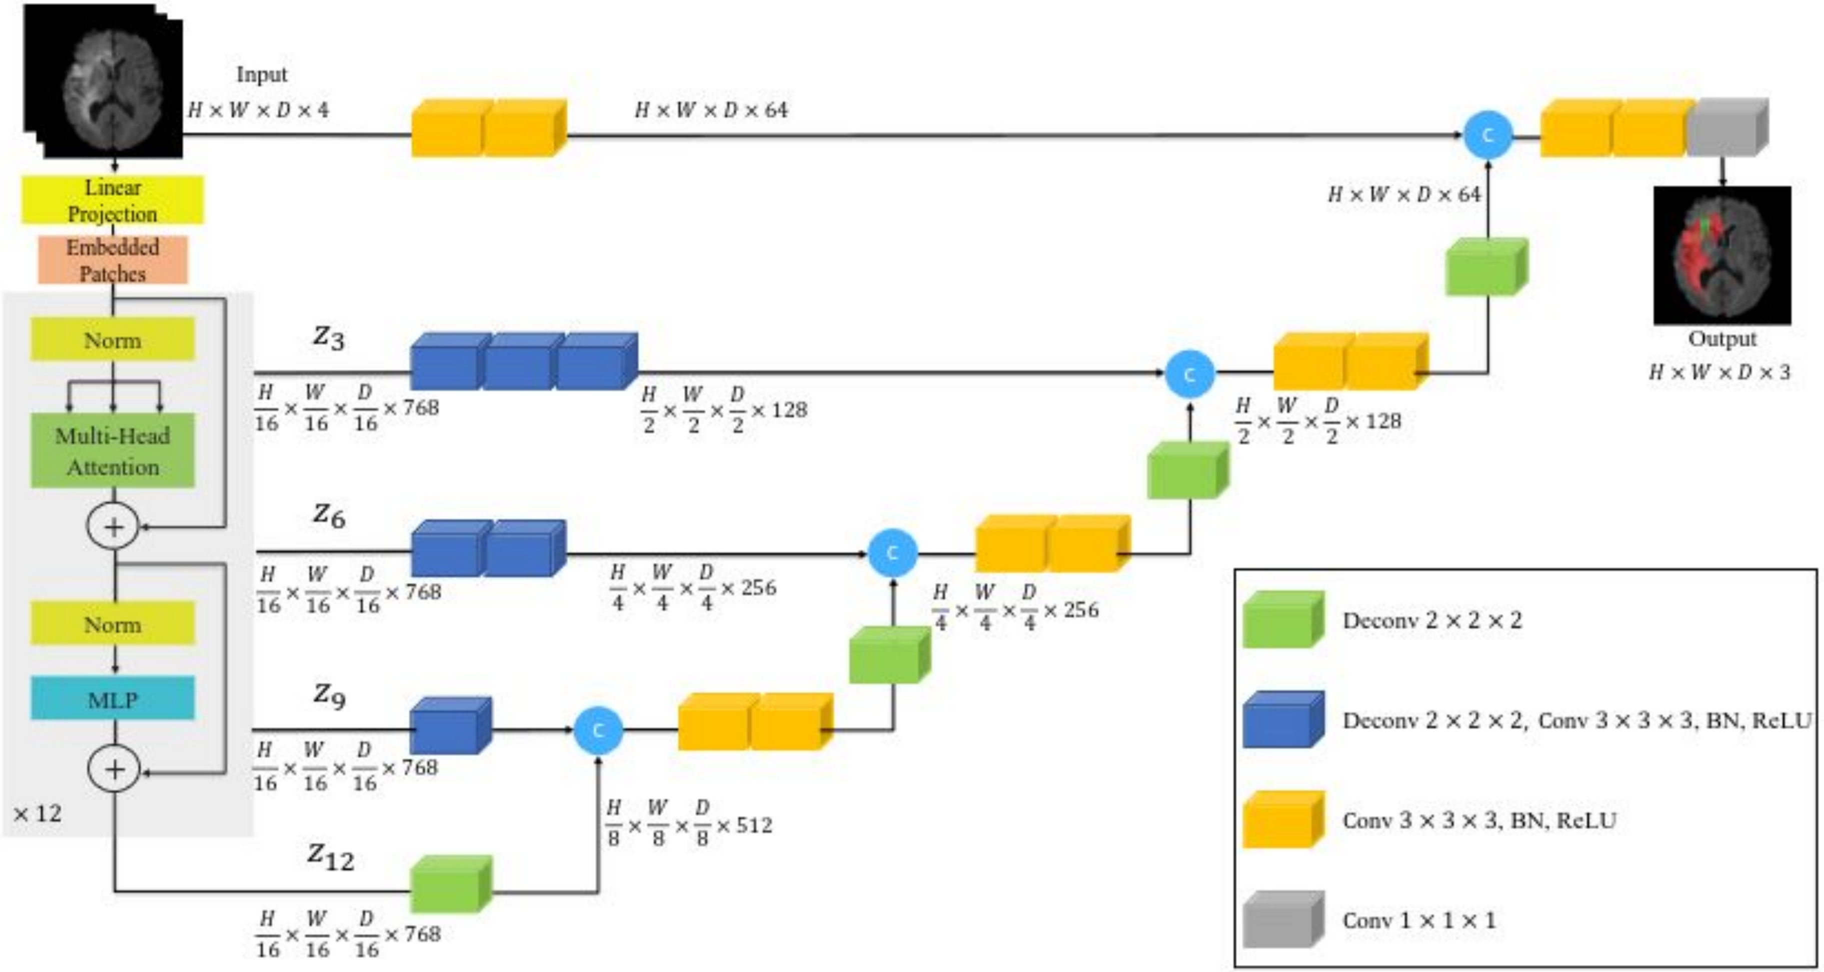

CNNs vs. Vision Transformer

- Main conceptual difference

- CNN: better or same for “small” data
  - every (annotated) dataset we have in microscopy is small!

- Transformers have advantages for large datasets and self-supervised learning

State-of-the-art DL / AI

Segment Anything Model and Applications in Microscopy

Segment Anything Model (SAM): Towards "Universal Segmentation"

https://segment-anything.com/

Model from Meta AI that can segment objects based on user input (left) / all objects in image (right)

SAM: How?

- Model architecture that can process user inputs (“prompts”)
- Very large training dataset, build iteratively with help of the model

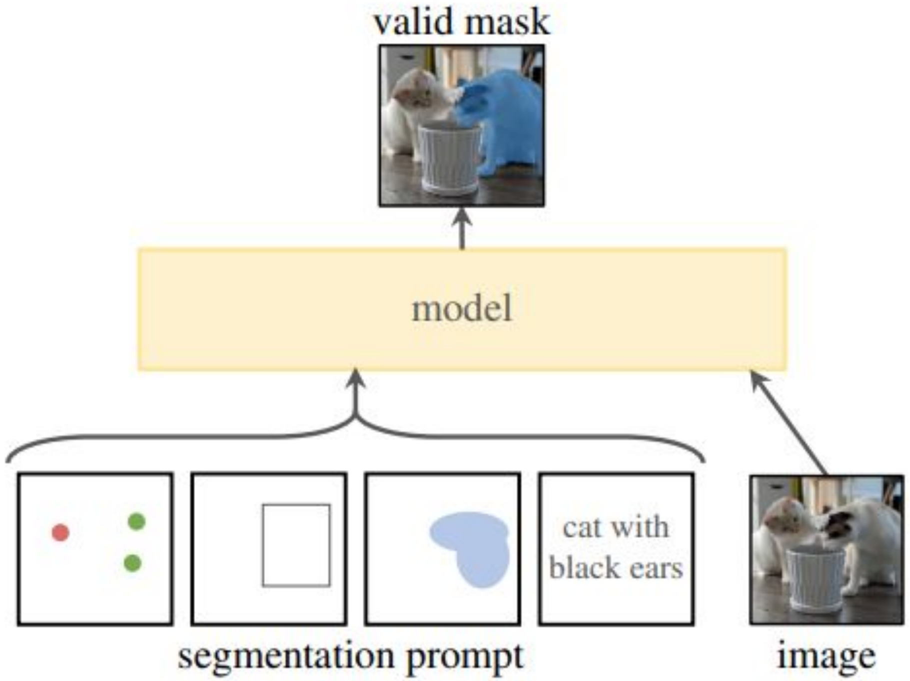

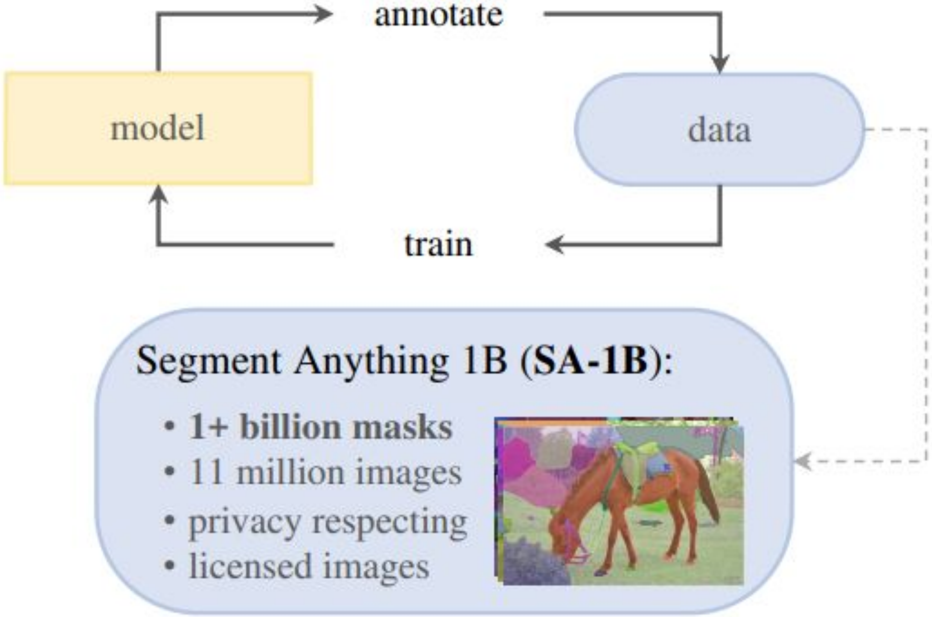

SAM: How?

- Model architecture that can process user inputs (“prompts”)
- Very large training dataset, build iteratively with help of the model

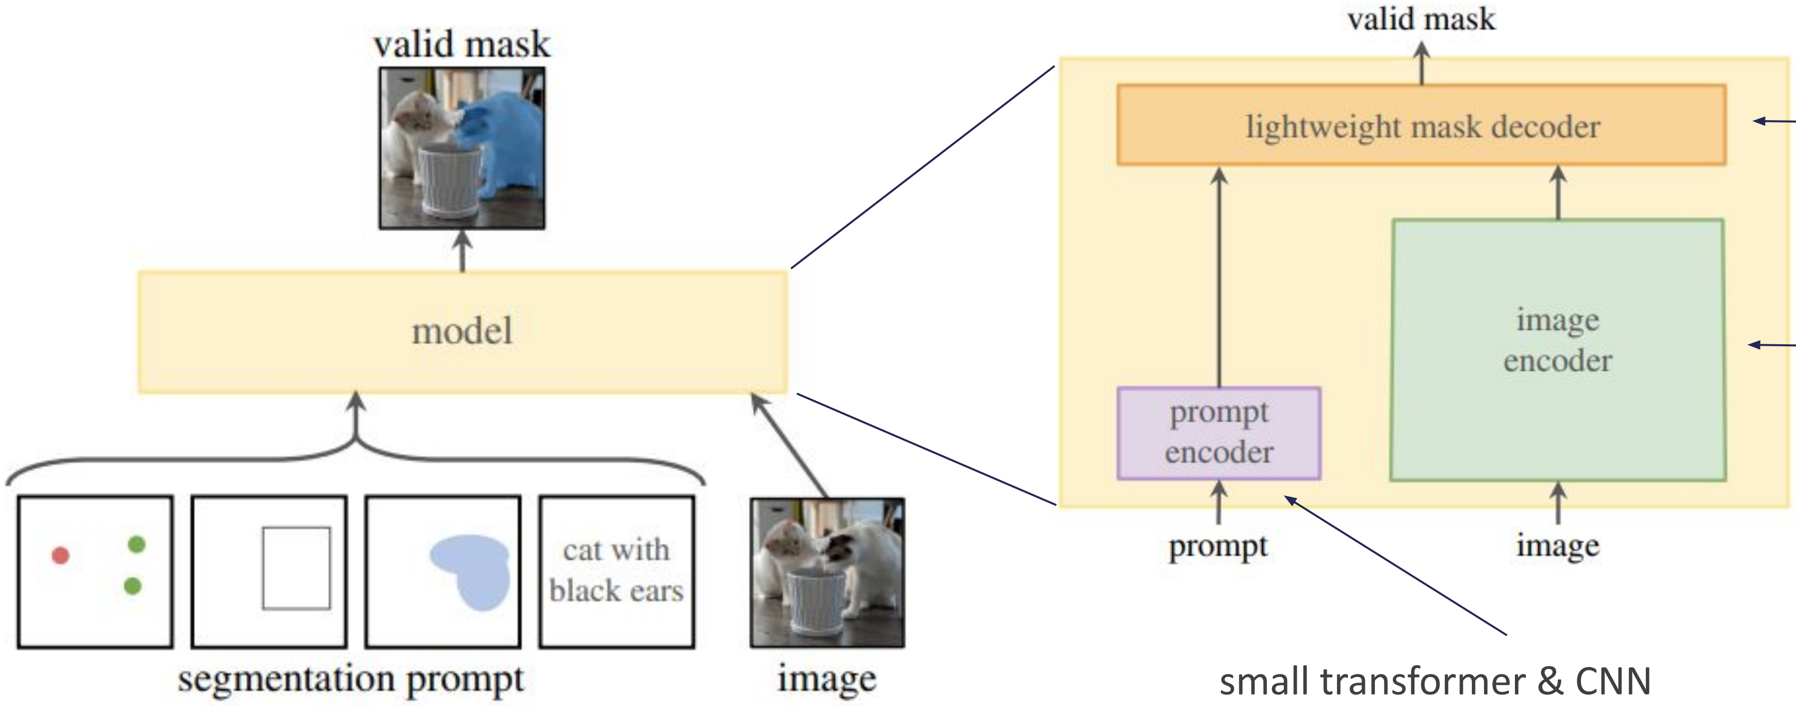

small transformer

vision transformer

small transformer & CNN

Segment Anything: Training iteration

Given image and ground-truth mask

- Compute image embeddings, sample positive point or box

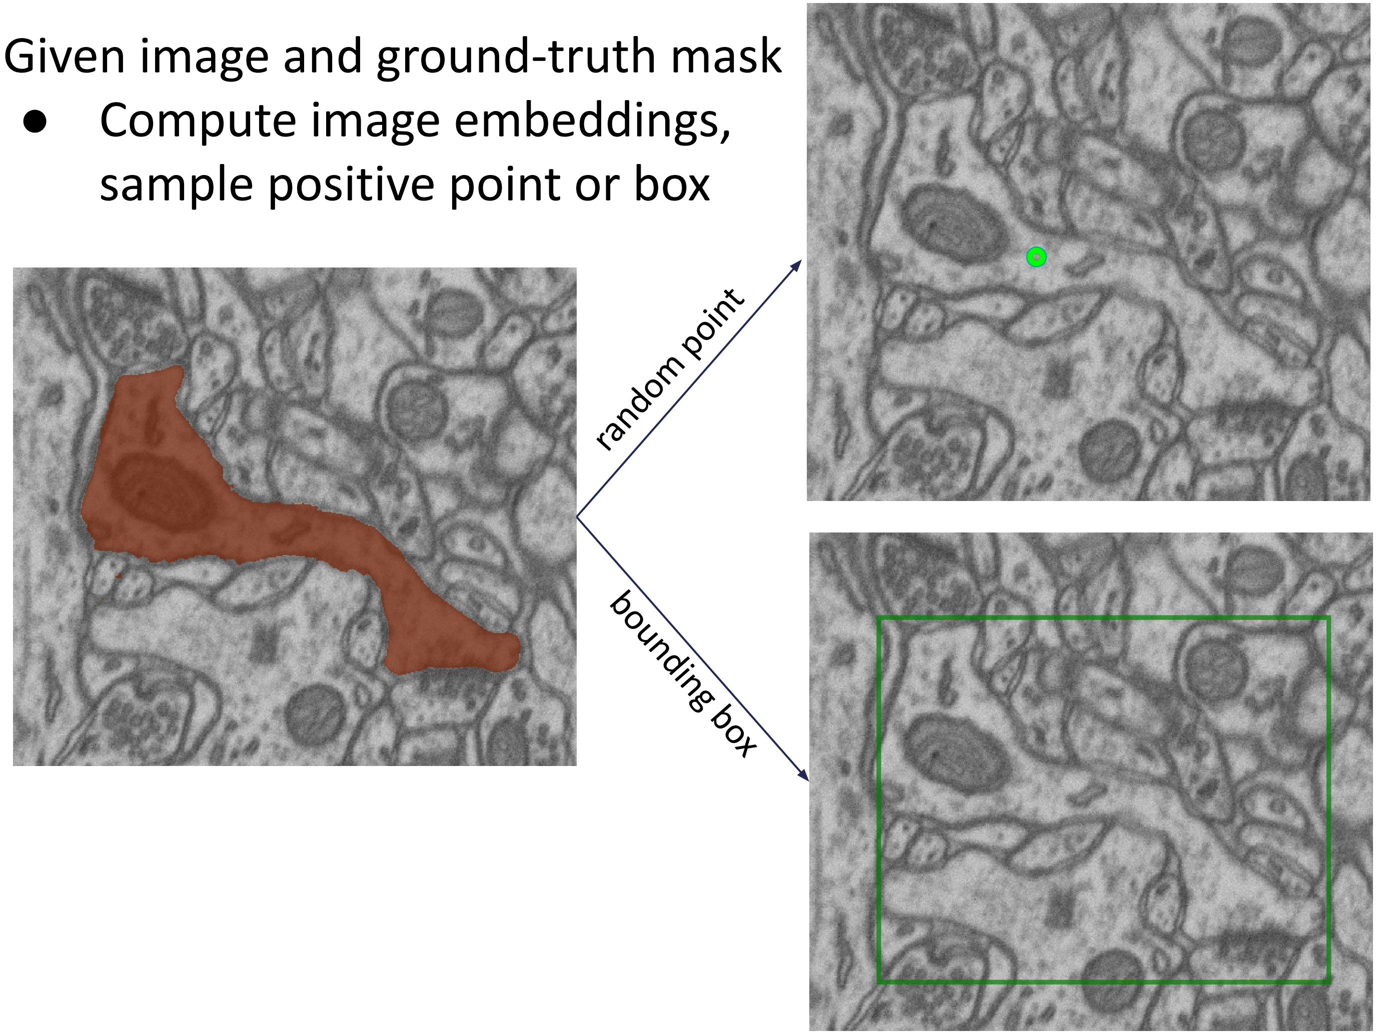

Segment Anything: Training iteration

Given image and ground-truth mask

- Compute image embeddings, sample positive point or box

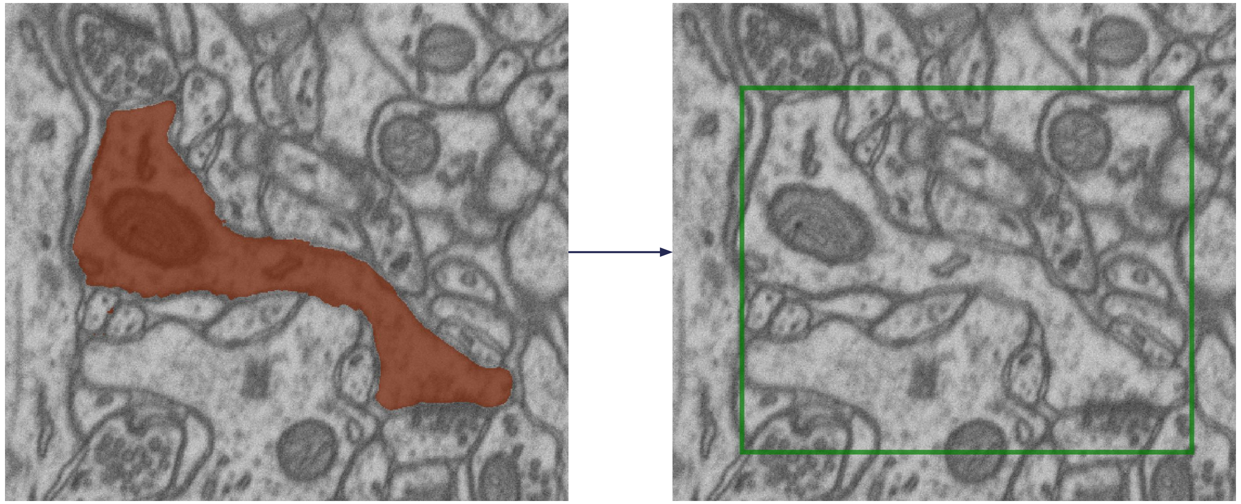

Segment Anything: Training iteration

Given image and ground-truth mask

- Compute image embeddings, sample positive point or box
- Run prediction, compute loss for object and IOU estimate

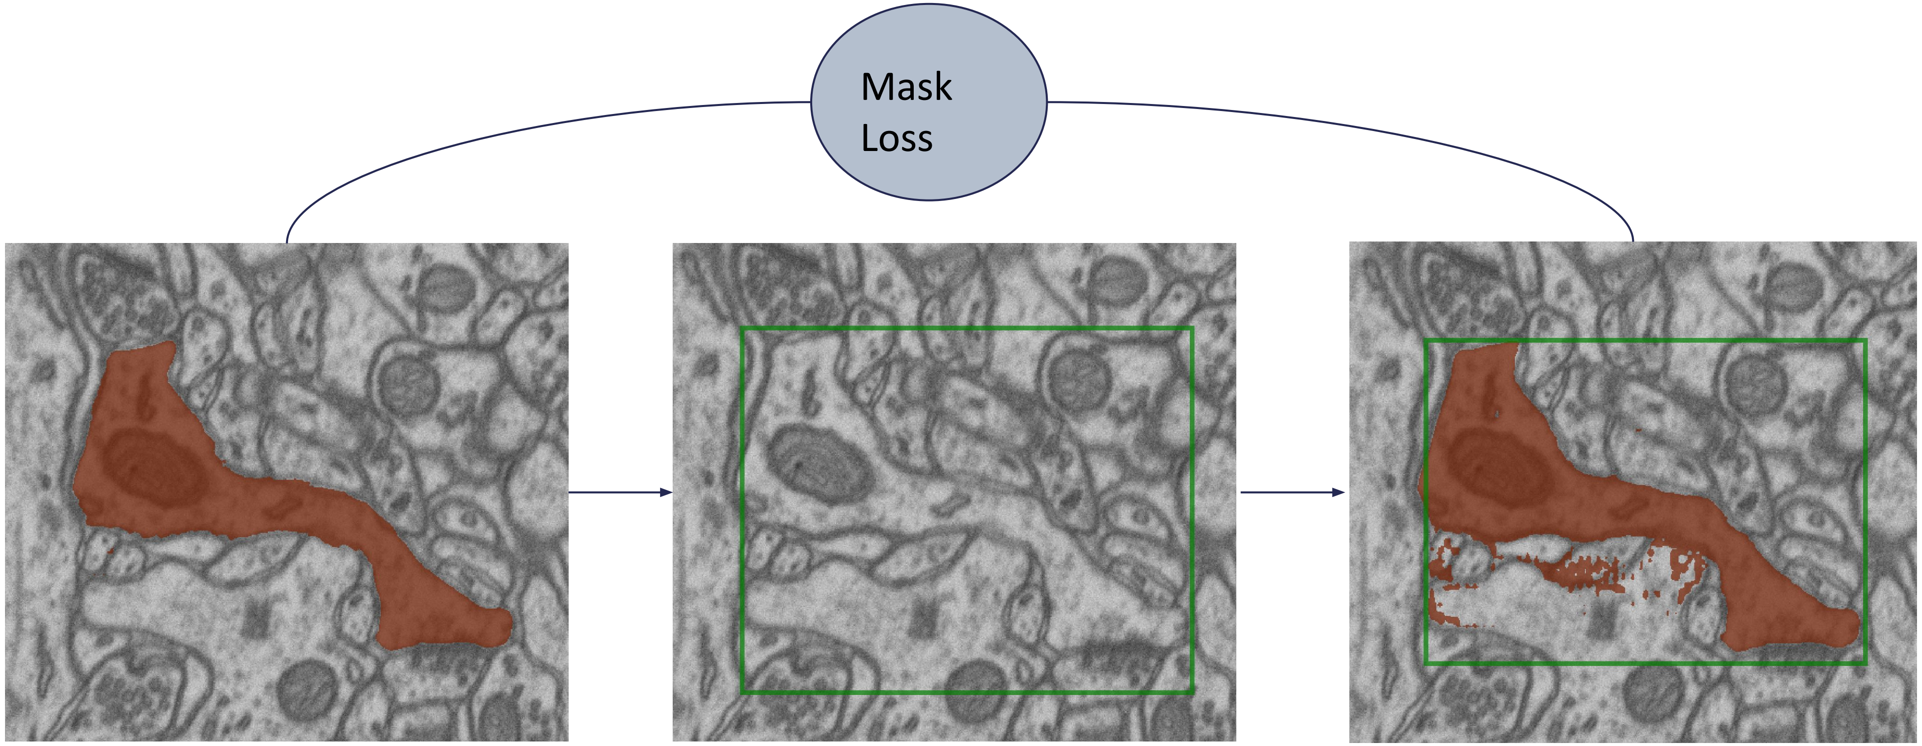

Segment Anything: Training iteration

Given image and ground-truth mask

- Compute image embeddings, sample positive point or box
- Run prediction, compute loss for object and IOU estimate
- Sample point prompts where prediction is wrong, rerun prediction with all prompts + mask

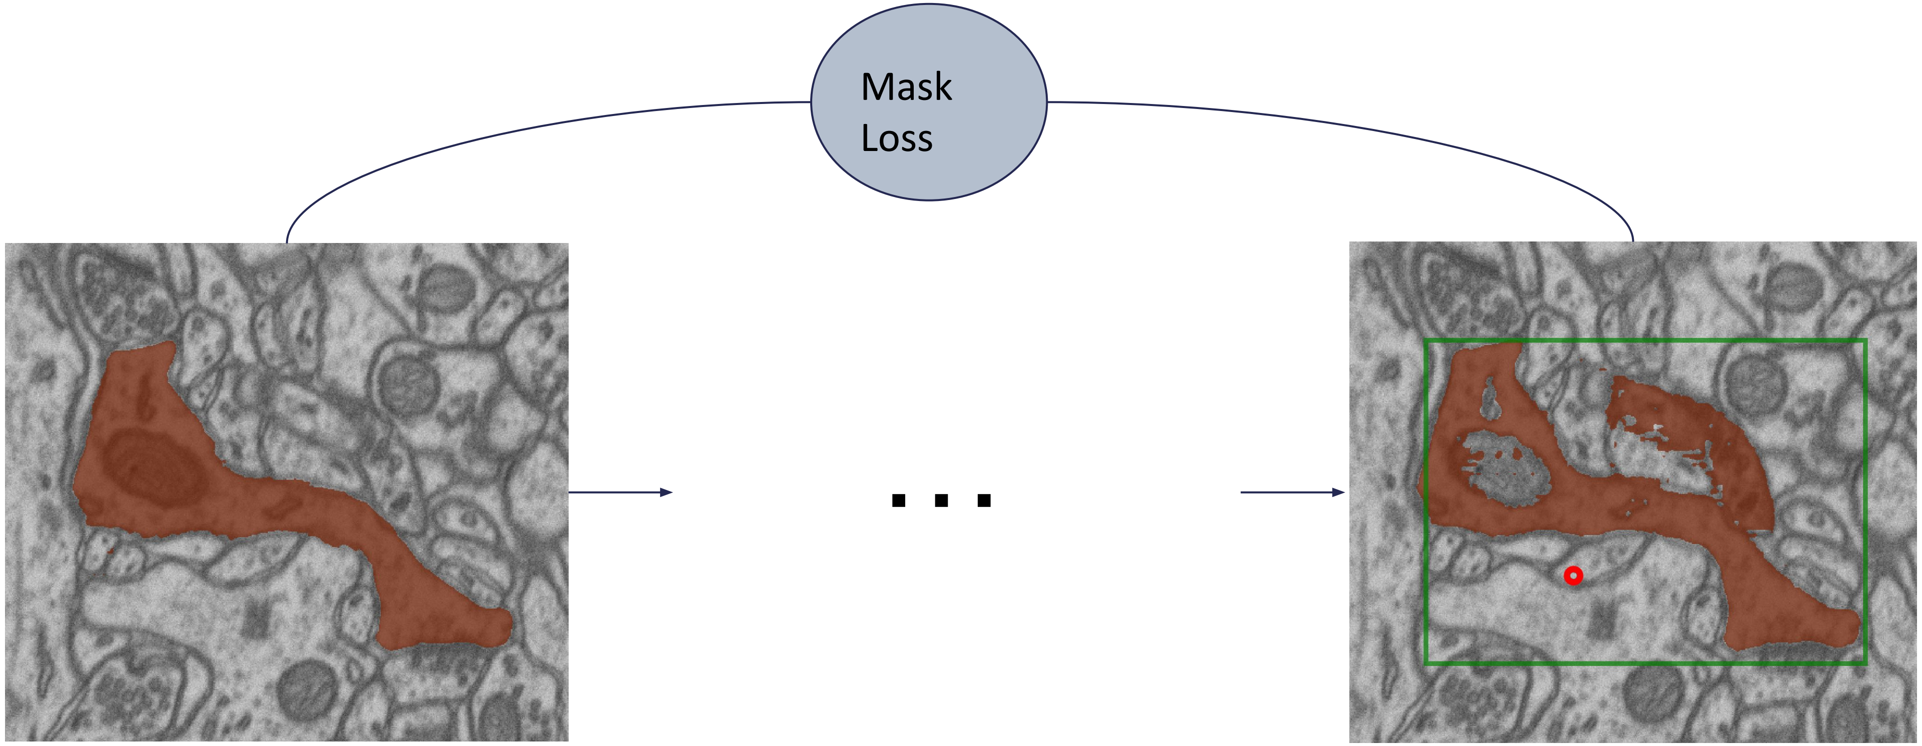

Segment Anything: Training iteration

Given image and ground-truth mask

- Compute image embeddings, sample positive point or box
- Run prediction, compute loss for object and IOU estimate
- Sample point prompts where prediction is wrong, rerun prediction with all prompts + mask
- Repeat

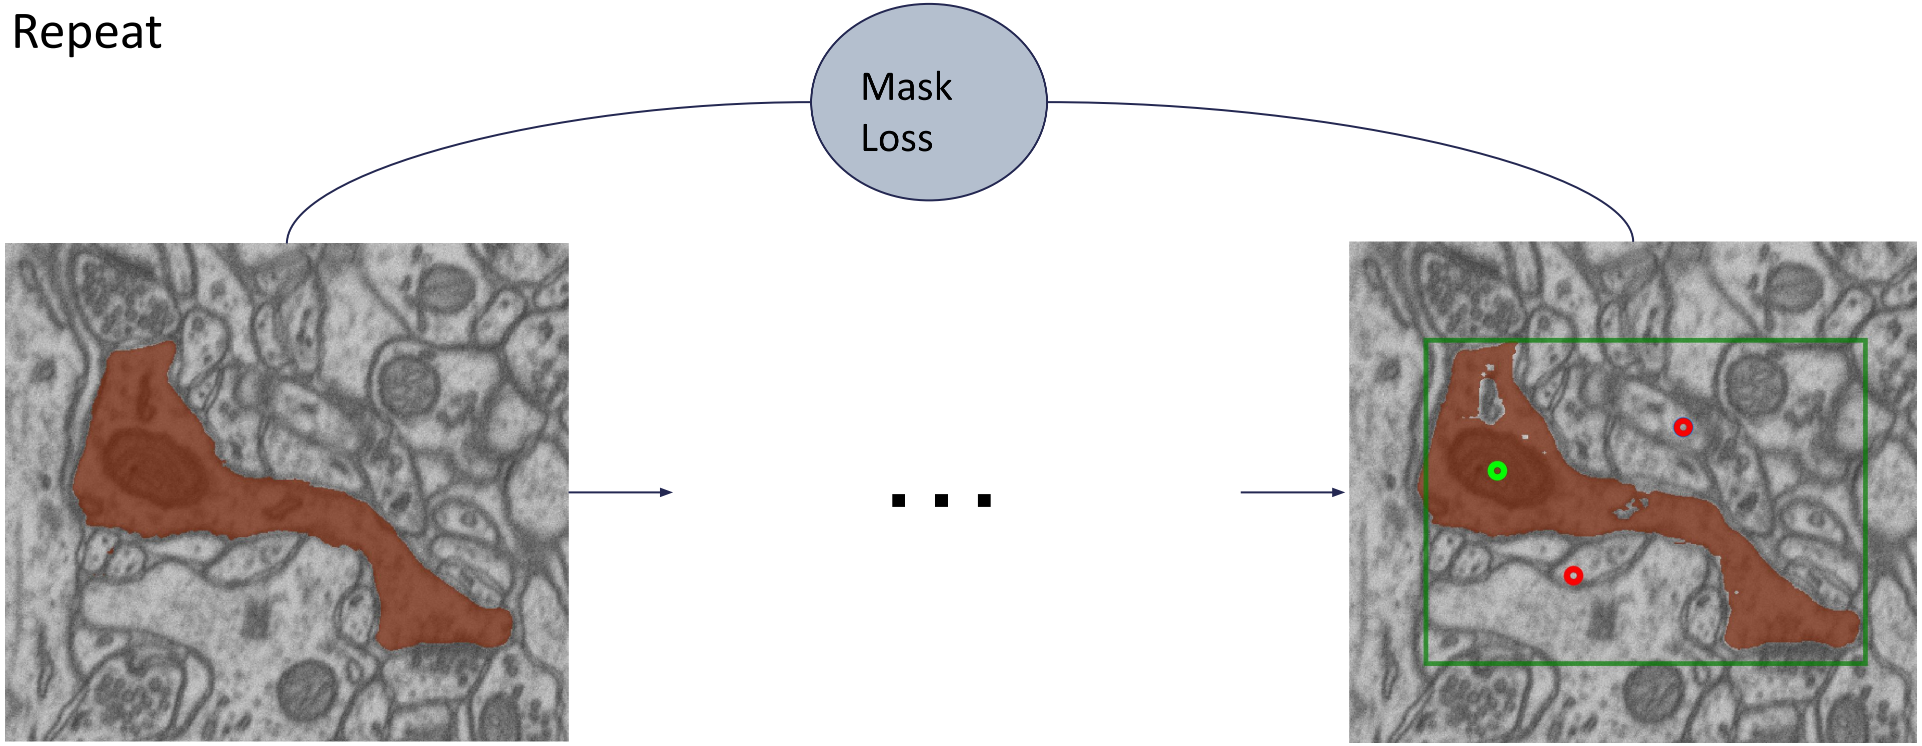

Segment Anything: Training iteration

Given image and ground-truth mask

- Compute image embeddings, sample positive point or box
- Run prediction, compute loss for object and IOU estimate
- Sample point prompts where prediction is wrong, rerun prediction with all prompts + mask
- Repeat

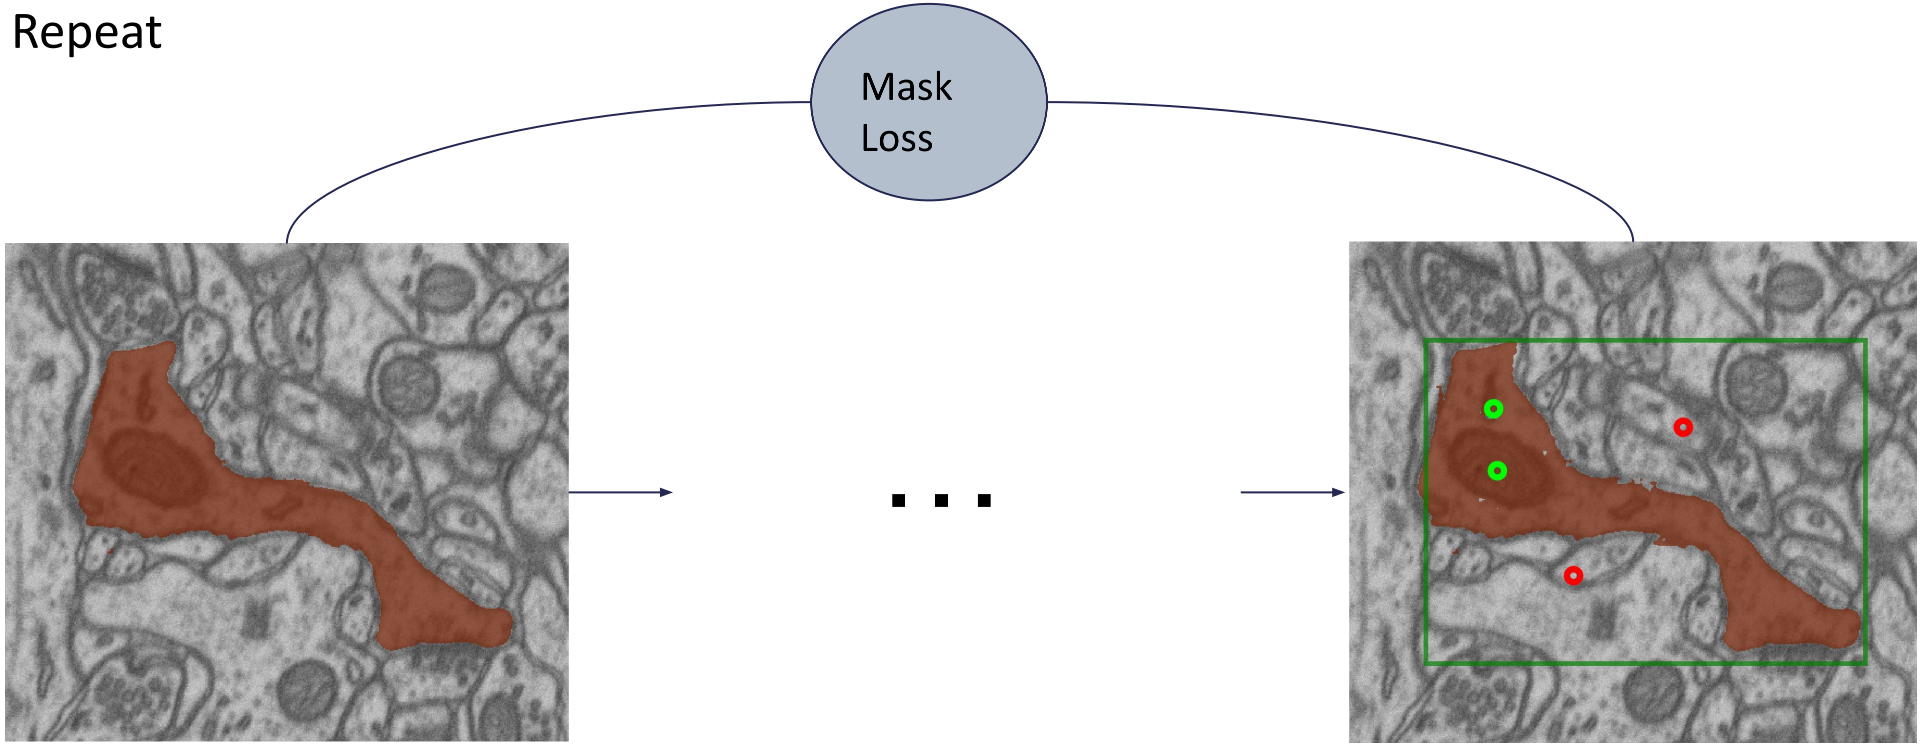

Mask
Loss

. . .

Segment Anything: Training iteration

Given image and ground-truth mask

- Compute image embeddings, sample positive point or box
- Run prediction, compute loss for object and IOU estimate
- Sample point prompts where prediction is wrong, rerun prediction with all prompts + mask
- Repeat
- Average losses, update weights

μSAM: Improving SAM for microscopy

μSAM: state-of-the-art method for microscopy segmentation based on SAM

- Re-trained on collection of public microscopy data with annotations
    - ○ Light Microscopy: 30k images; >3 million annotated cells and nuclei
    - ○ Electron Microscopy: 5k images; >100k annotated mitochondria and nuclei
- New segmentation (U-Net-like) decoder for improved automatic segmentation

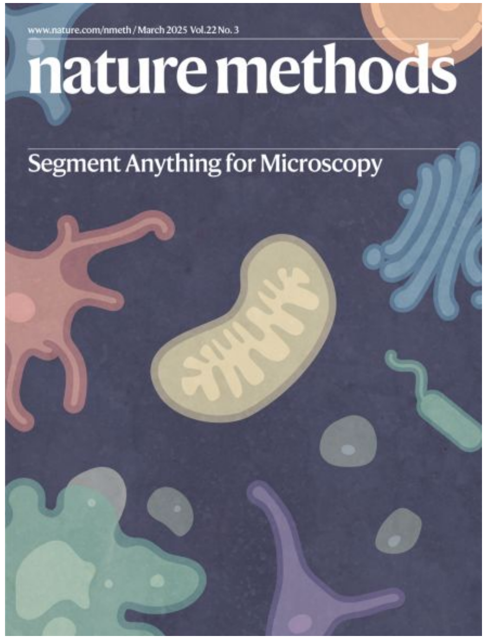

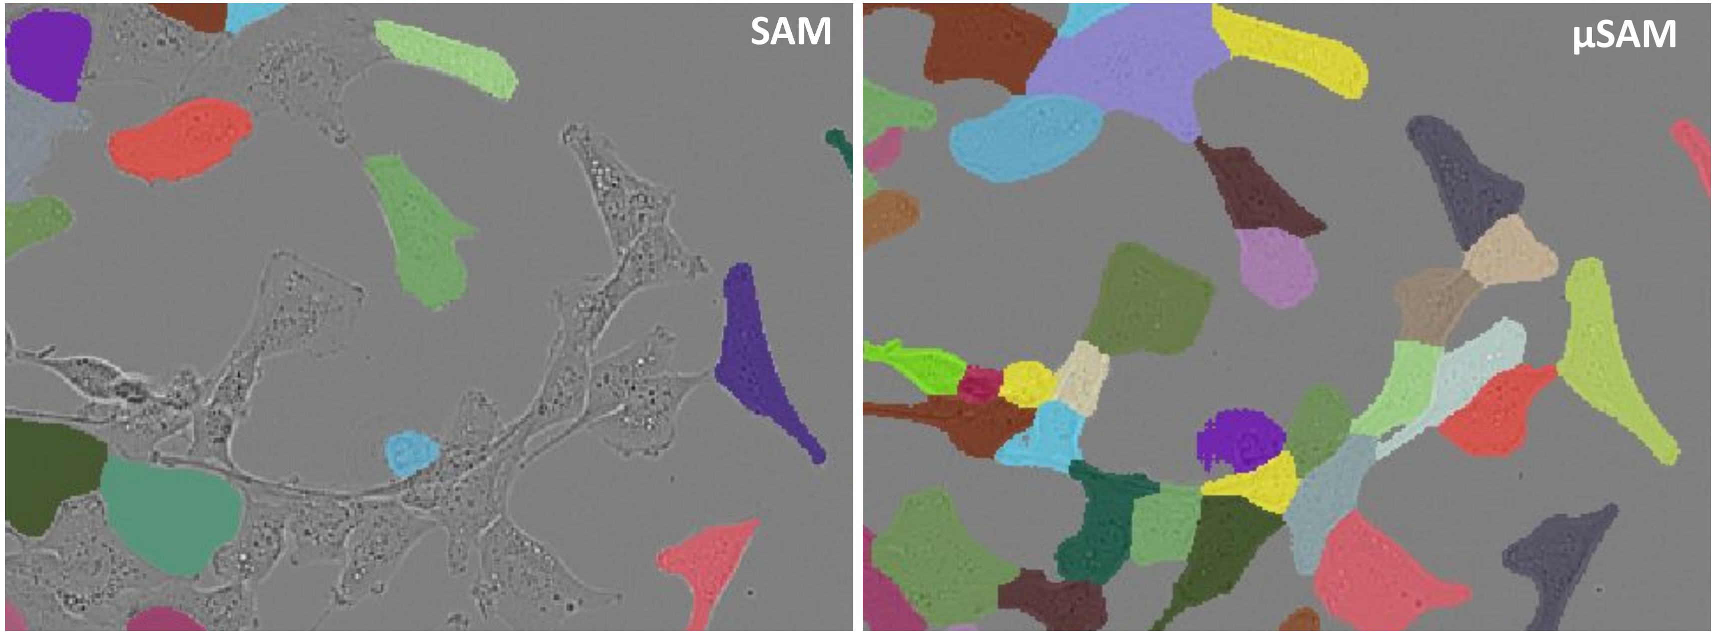

Archit, …, **Pape**. *Nature Methods* 2025

Interactive Segmentation with µSAM

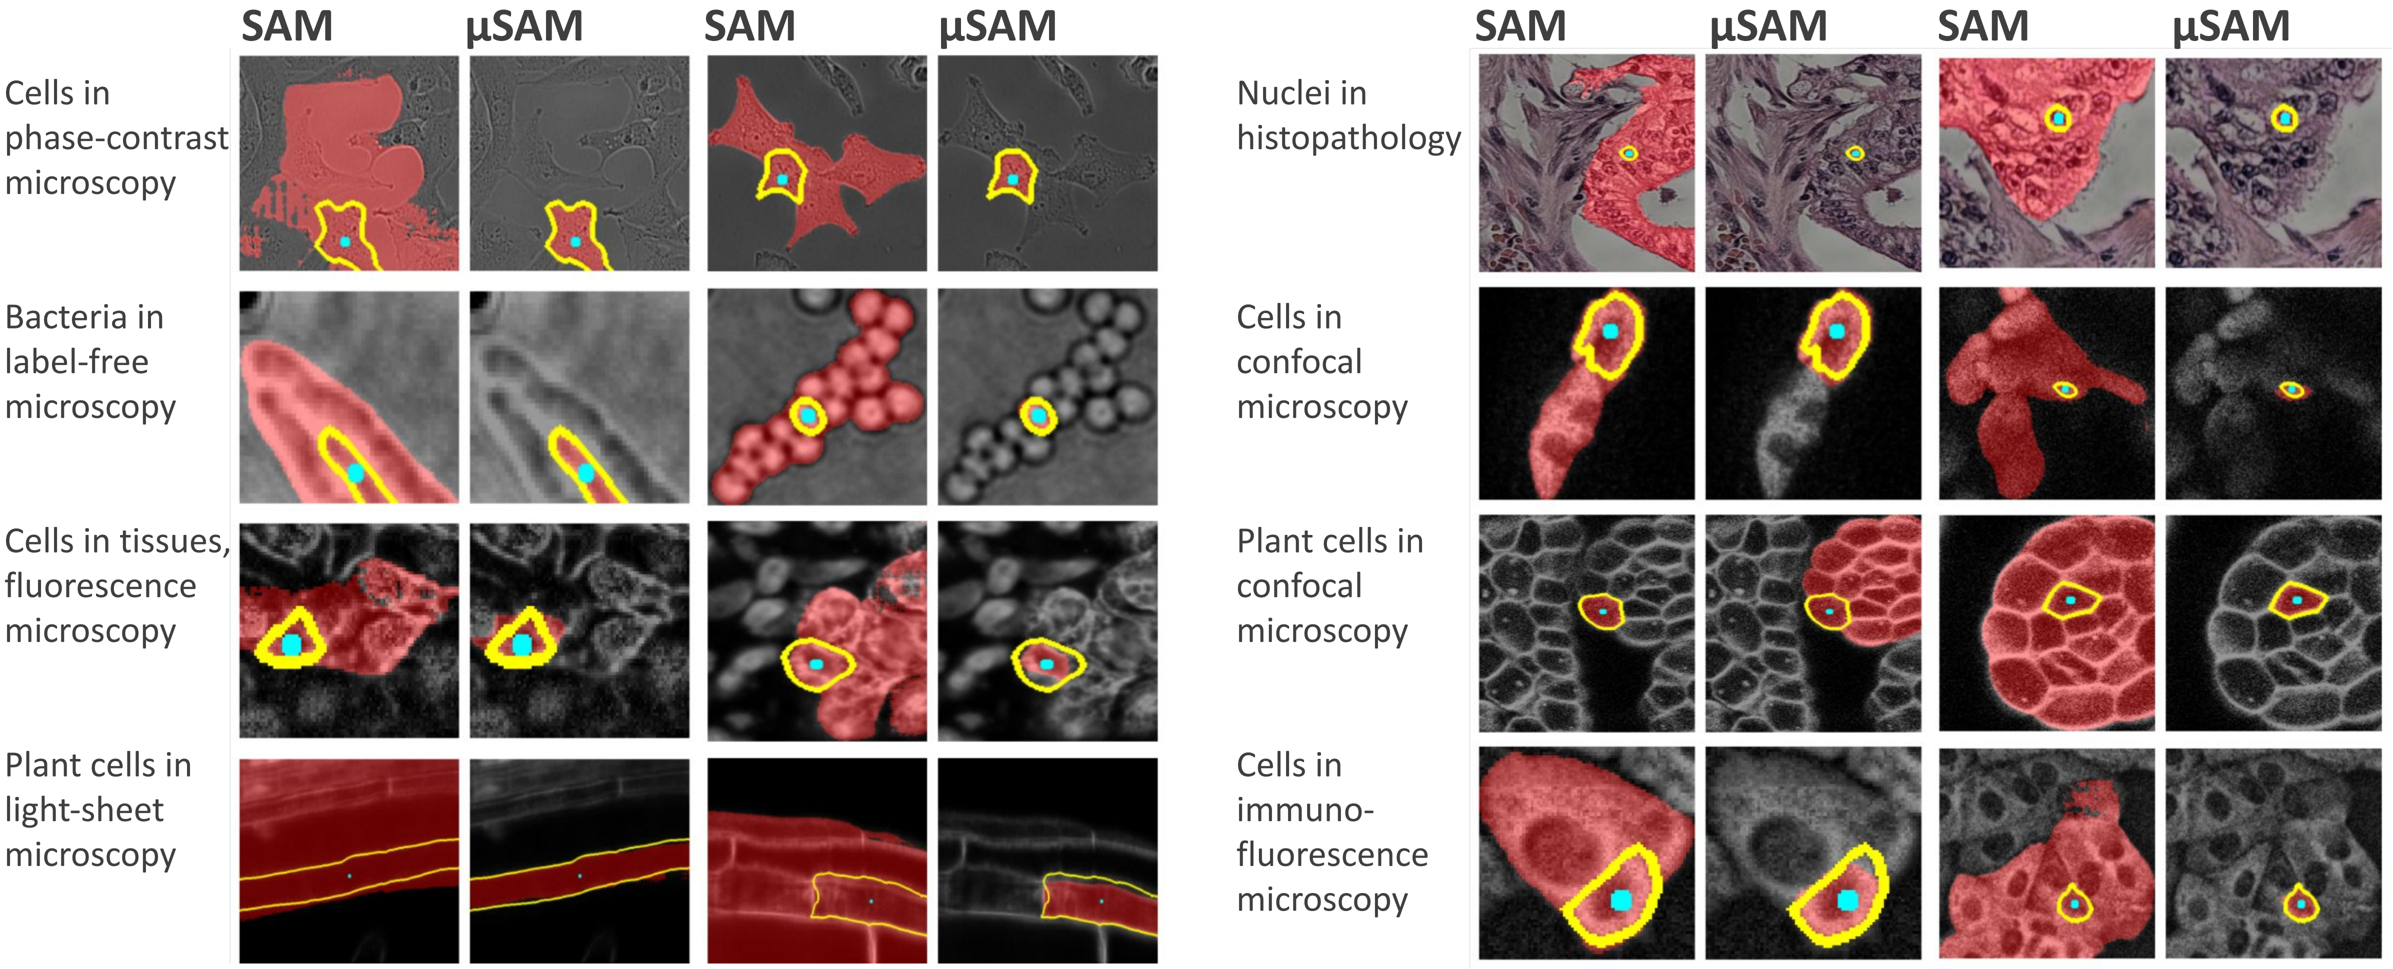

Automatic Segmentation with µSAM

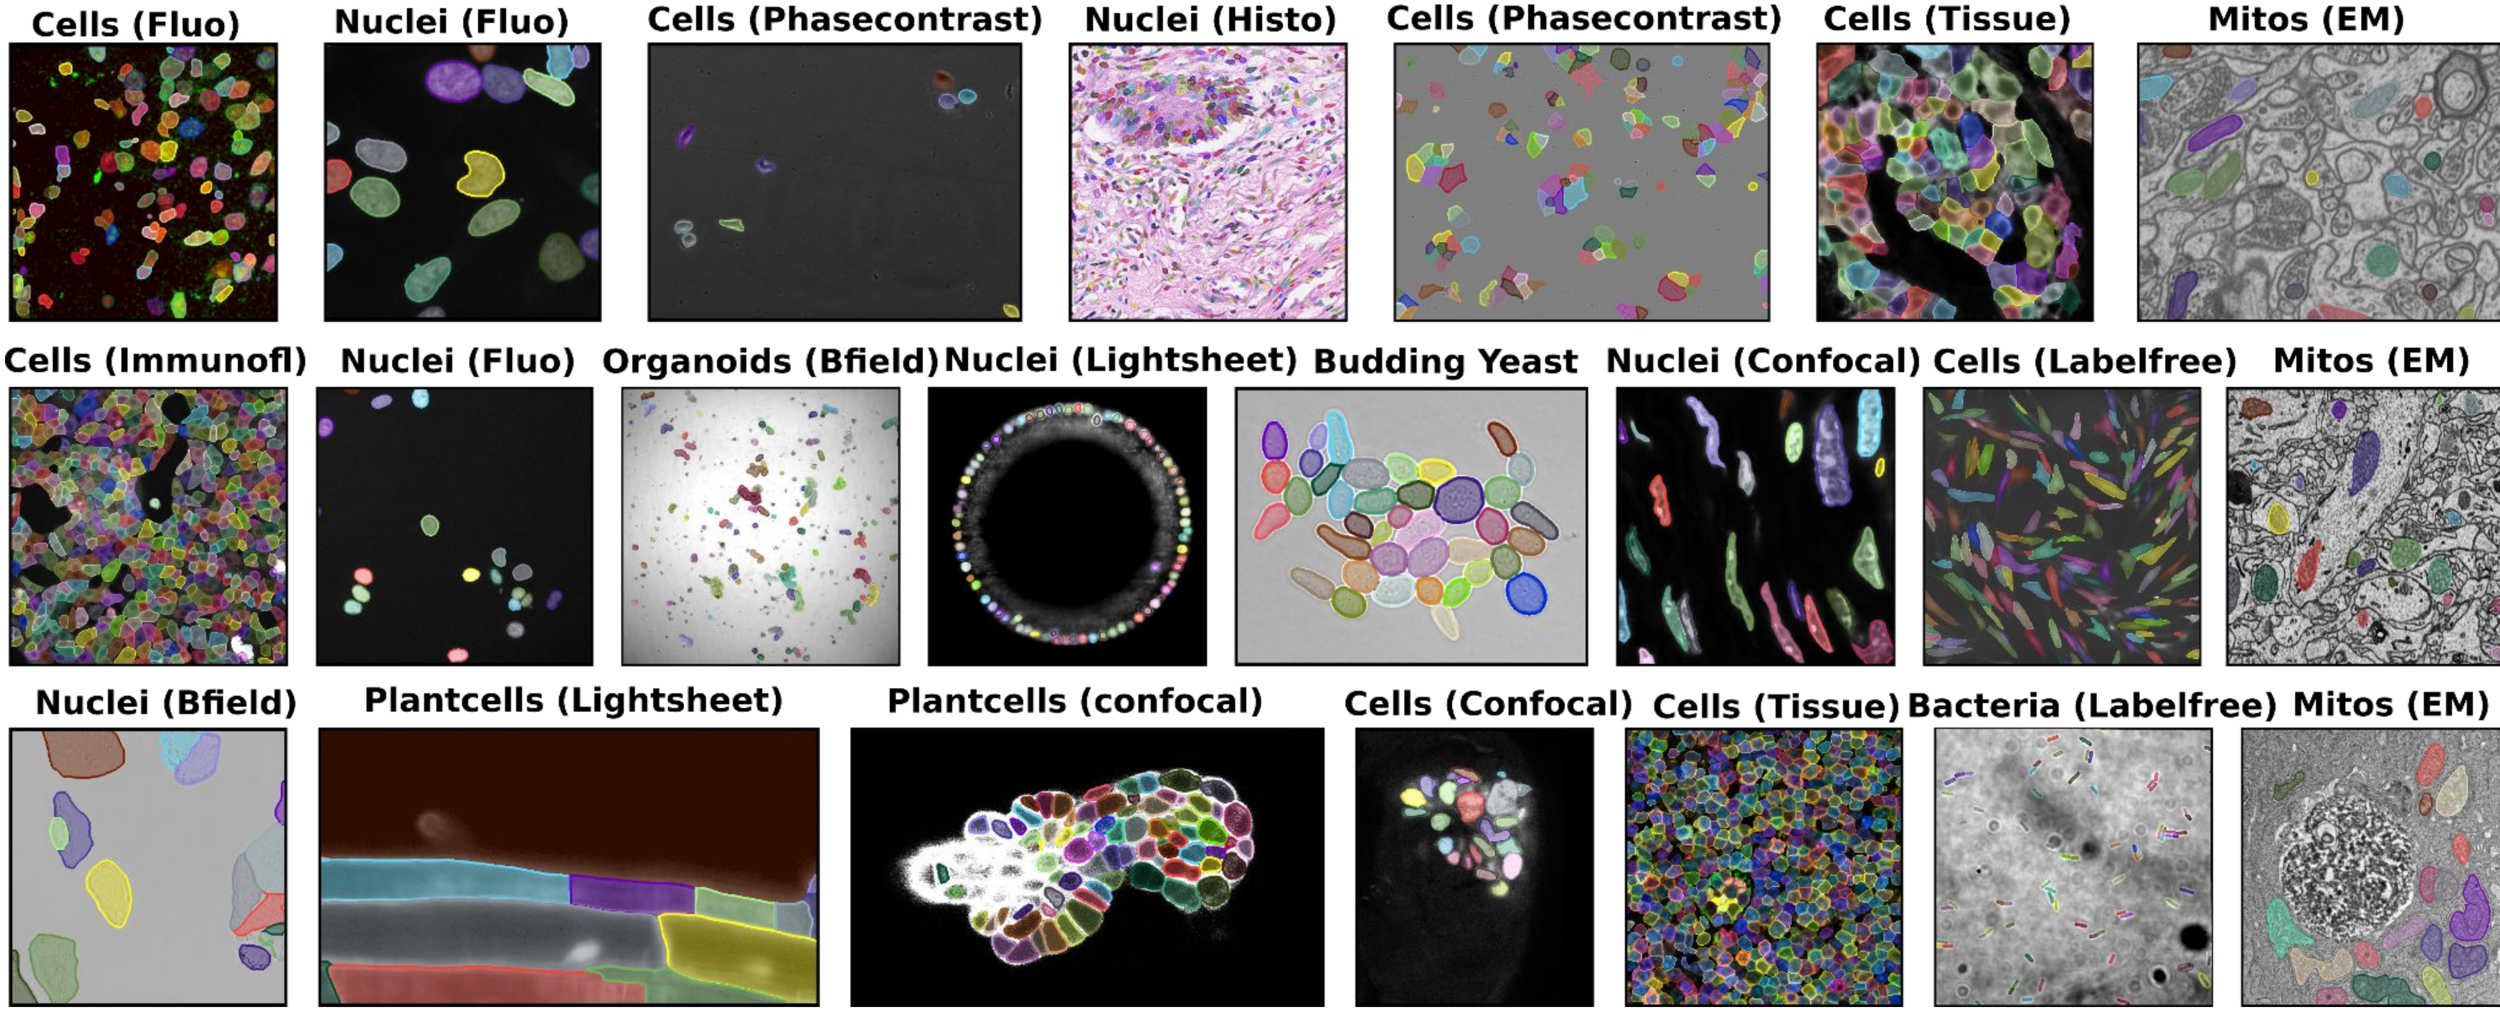

µSAM: user-friendly tool for microscopy segmentation

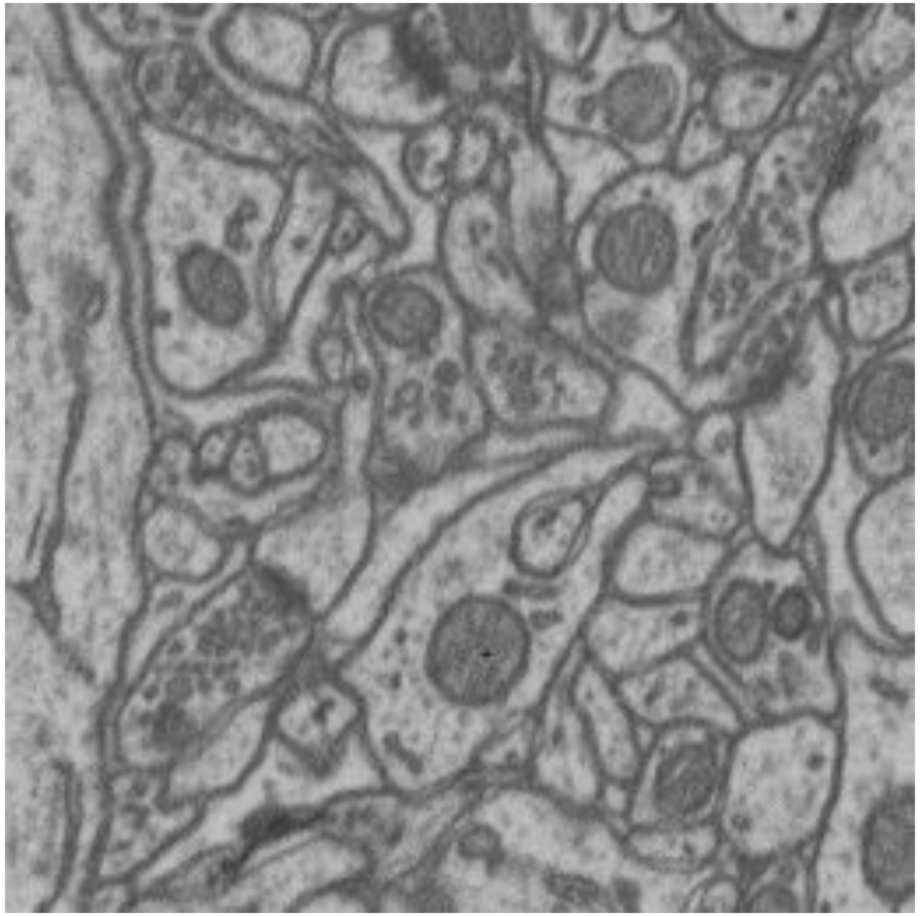

Available at https://computational-cell-analytics.github.io/micro-sam/

µSAM: user-friendly tool for microscopy segmentation

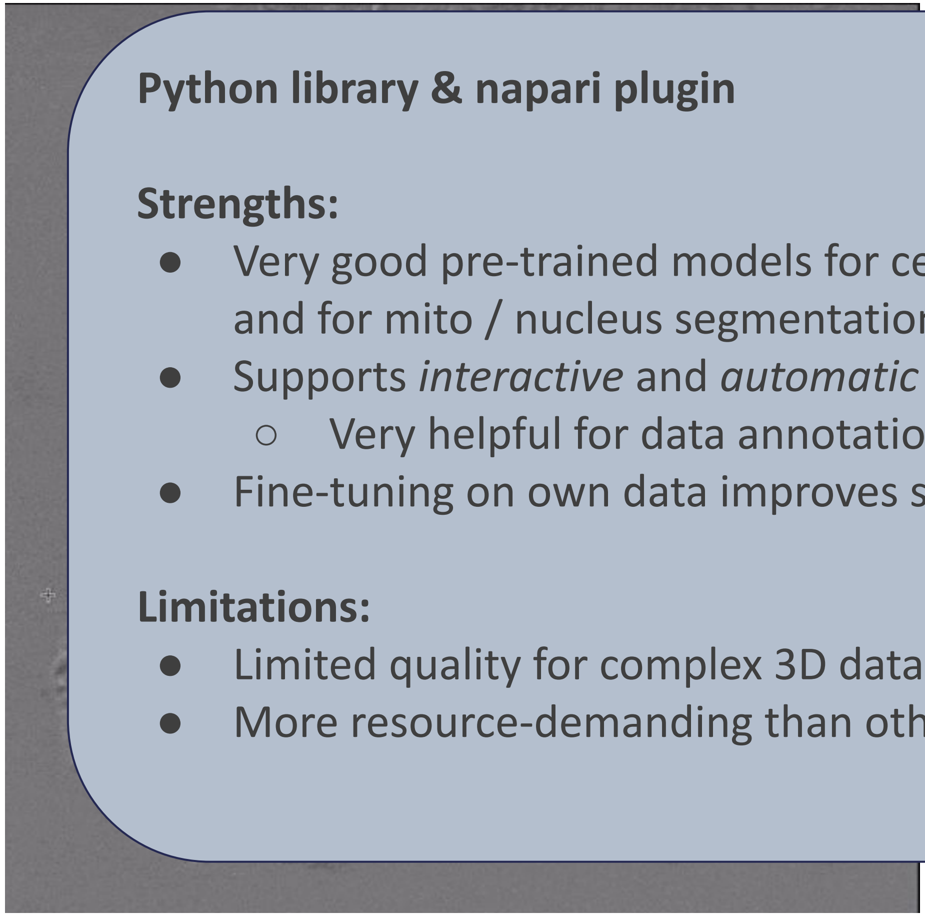

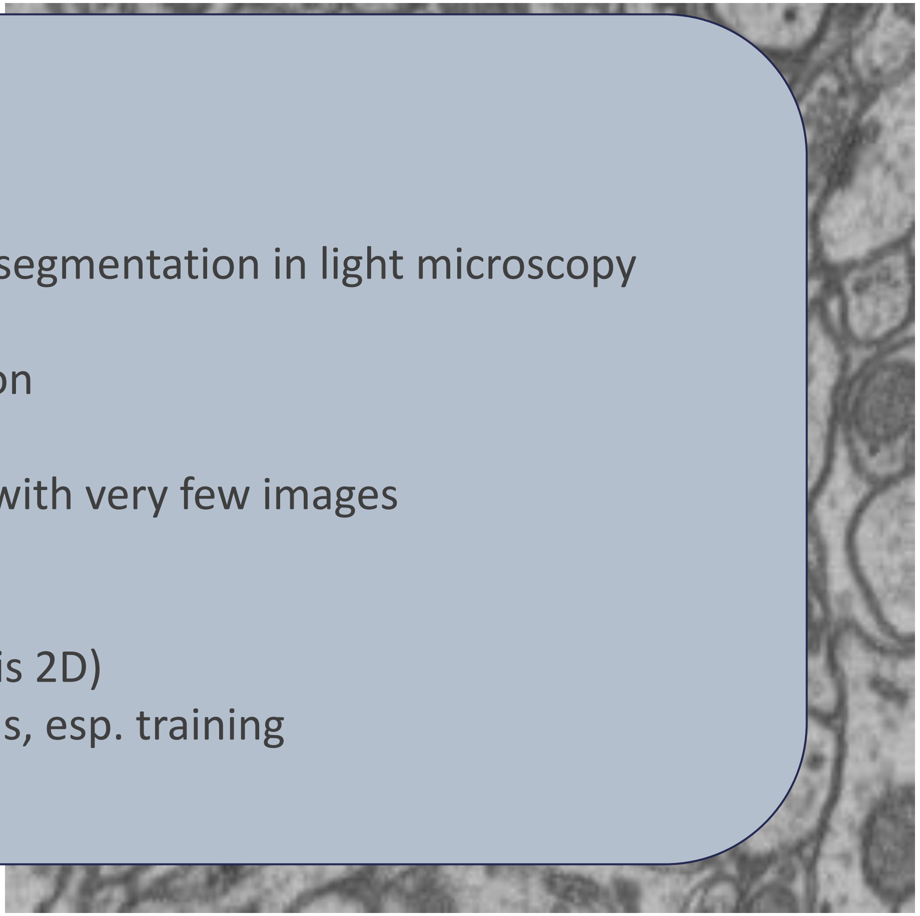

Python library & napari plugin

Strengths:

- Very good pre-trained models for cell / nucleus segmentation in light microscopy and for mito / nucleus segmentation in EM
- Supports *interactive* and *automatic* segmentation
  - Very helpful for data annotation
- Fine-tuning on own data improves significantly with very few images

Limitations:

- Limited quality for complex 3D data (SAM itself is 2D)
- More resource-demanding than other DL models, esp. training

Available at https://computational-cell-analytics.github.io/micro-sam/

CellPose SAM

https://www.cellpose.org/

Newest version of CellPose (“CellPose 4”) also finetunes SAM for cell segmentation.

Integrated in CellPose library / tool

Strengths:

- Very good pre-trained model for cell segmentation

Limitations:

- No interactive segmentation
- More resource-demanding than other DL models, esp. training

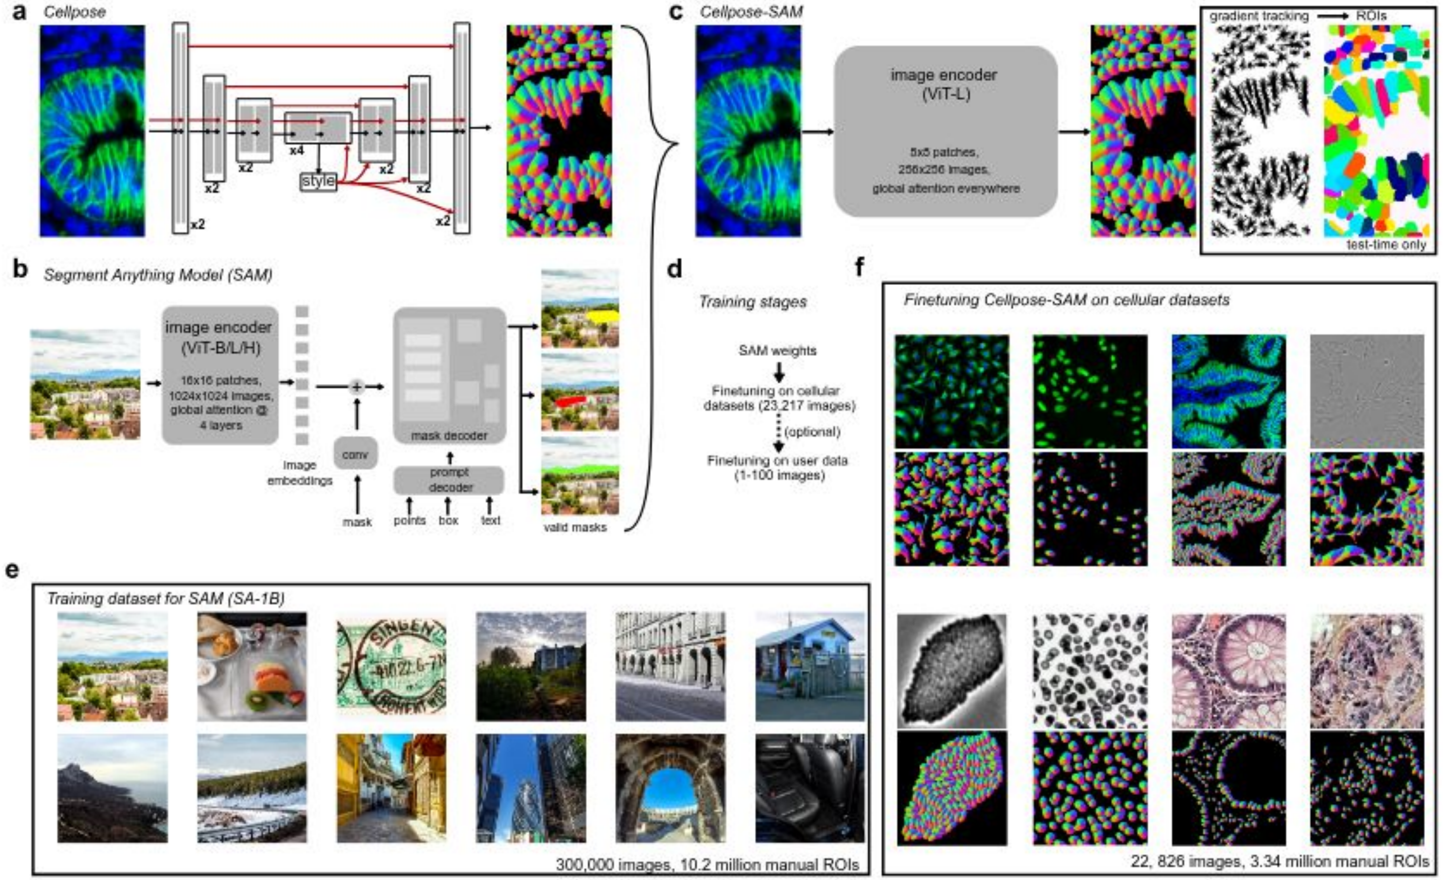

Outlook SAM2: Interactive video segmentation

Extension of SAM: interactive video segmentation
More efficient model for multi-dimensional data

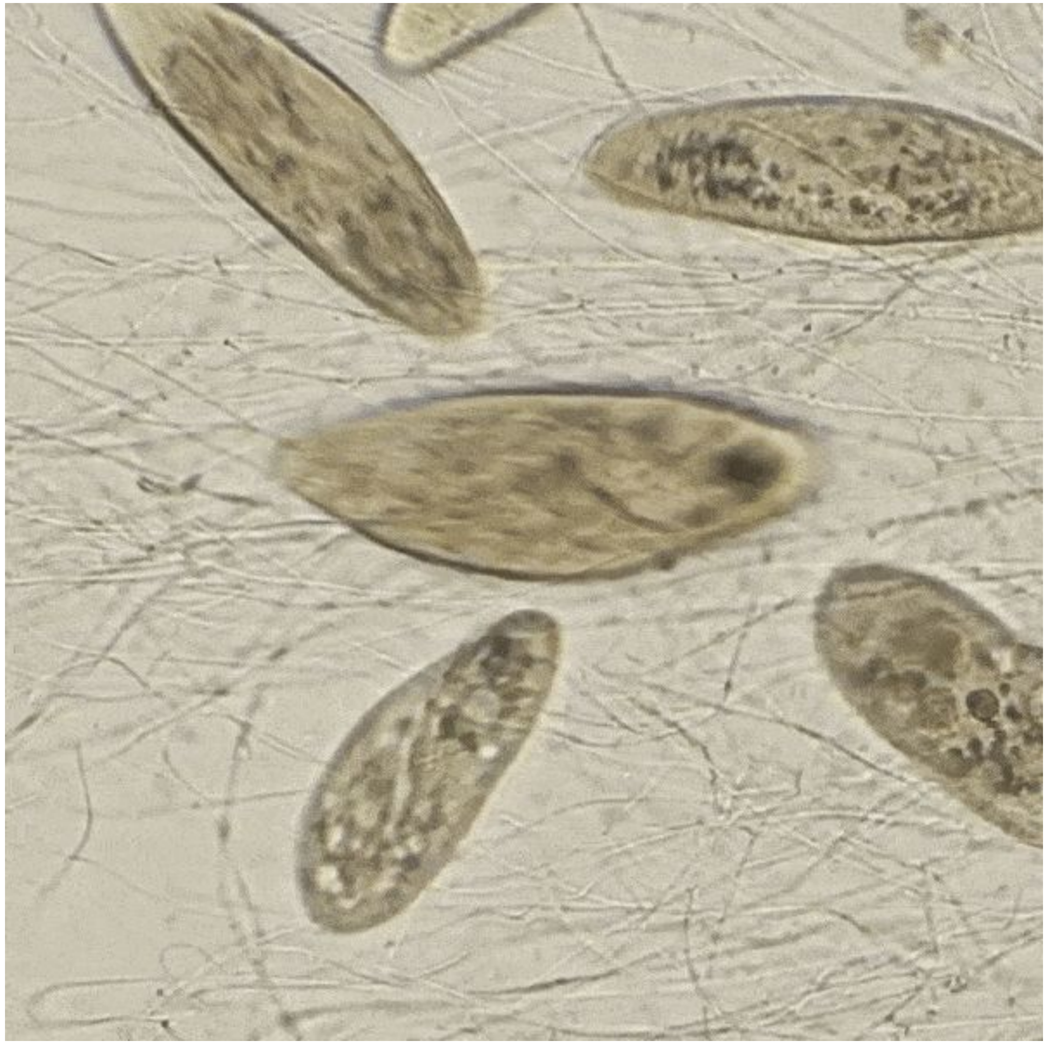

Outlook SAM2: Interactive video segmentation

Extension of SAM: interactive video segmentation
More efficient model for multi-dimensional data

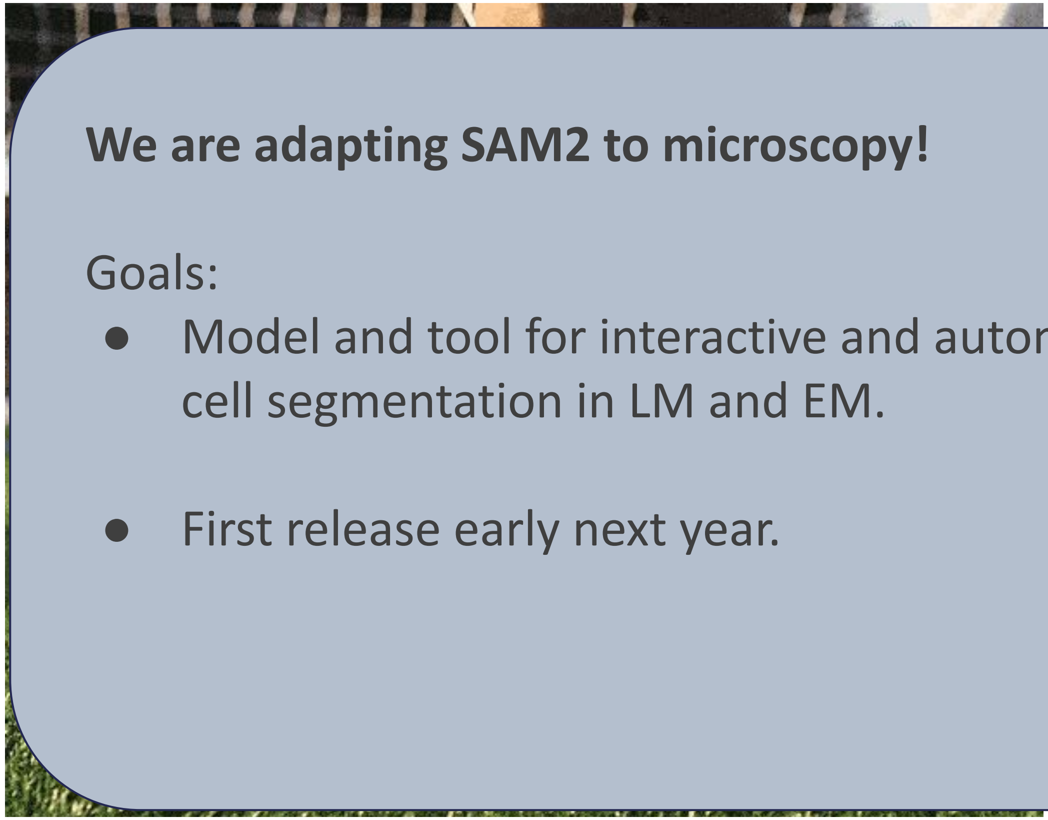

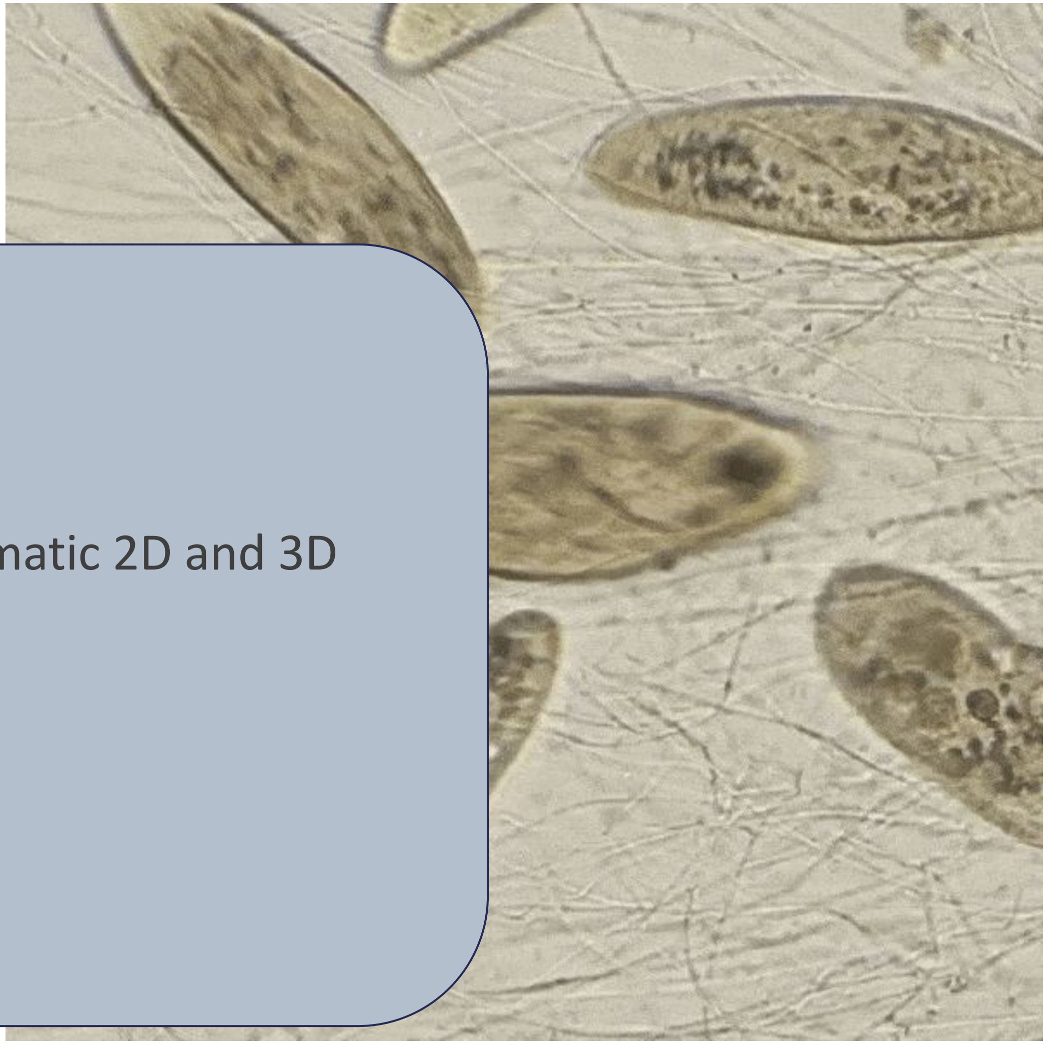

We are adapting SAM2 to microscopy!

Goals:

- Model and tool for interactive and automatic 2D and 3D cell segmentation in LM and EM.
- First release early next year.

Data Annotation

Segmentation annotation in napari

Live demo!

Segmentation annotation: Caveats

Annotation has to be **dense**: every cell / object in the images / image patches used for training has to be annotated!

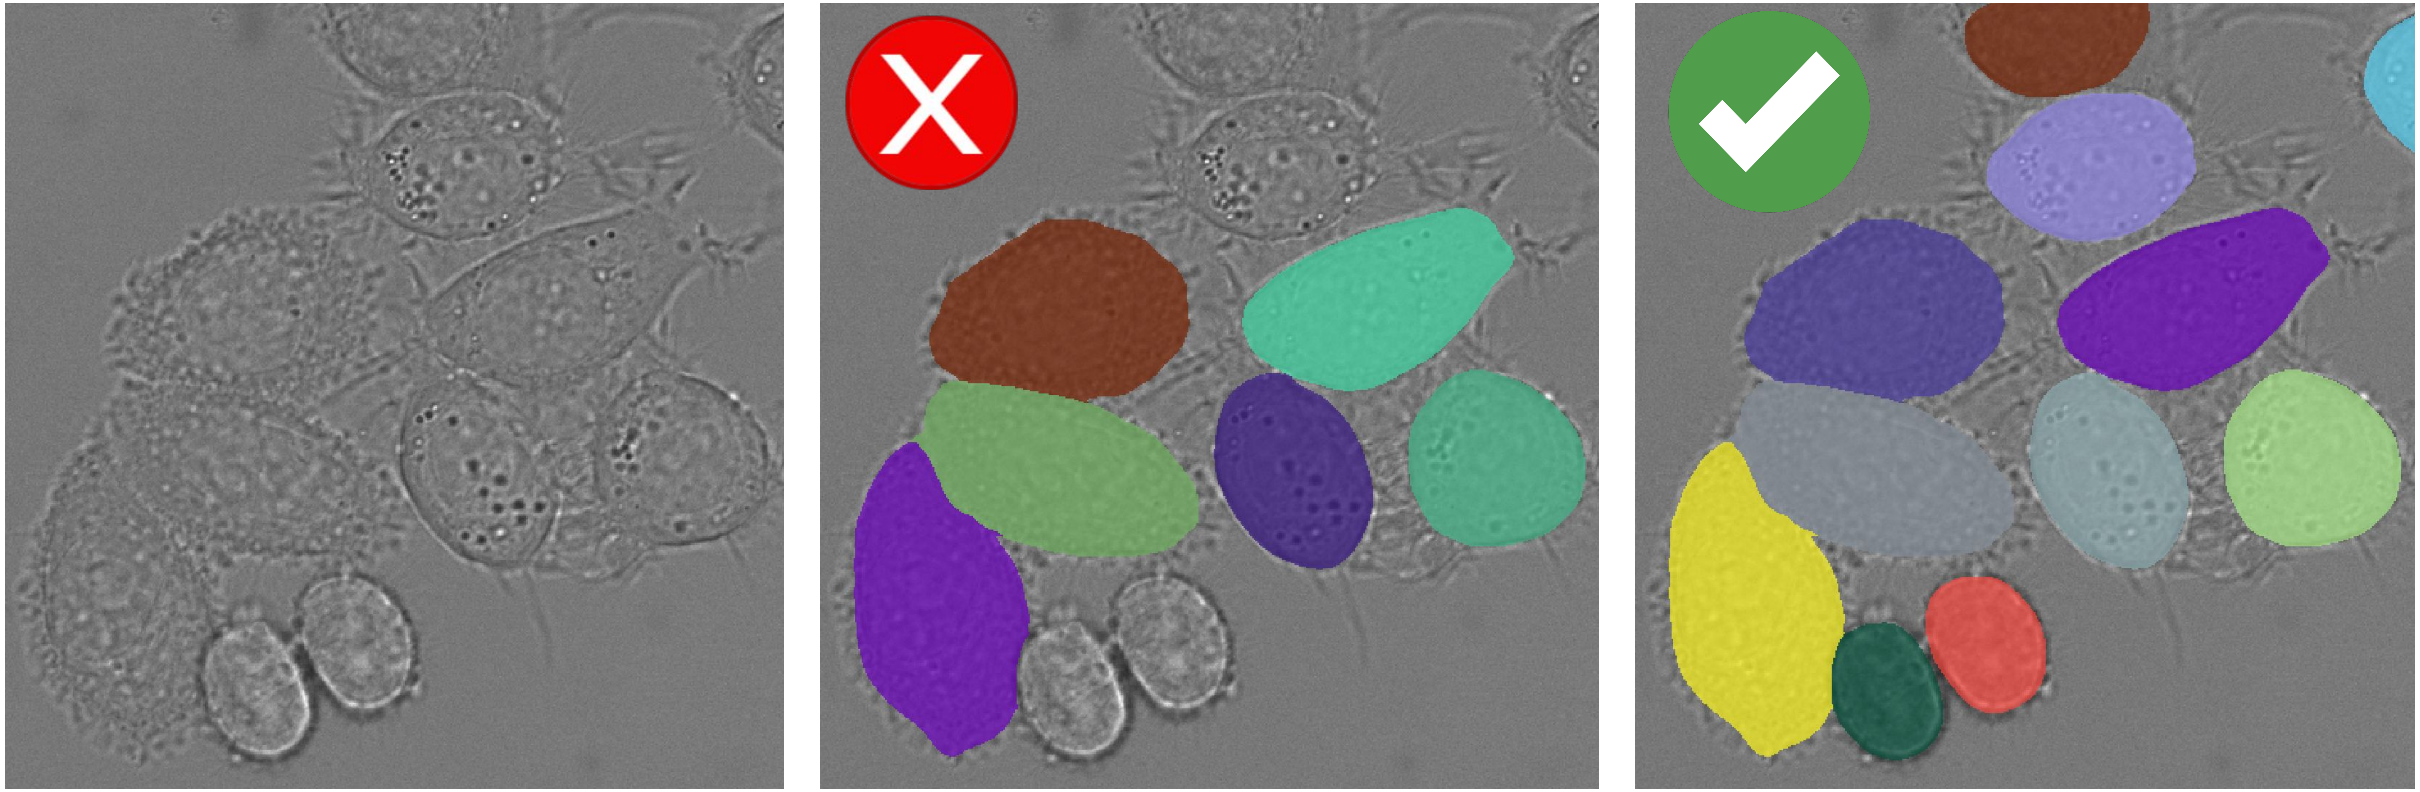

Segmentation annotation: Caveats

Annotation has to be **dense**: every cell / object in the images / image patches used for training has to be annotated!

Very tedious! Let’s get some tools to help us!

Segmentation annotation with µSAM

Live demo!

Interactive annotation (+ auto segmentation) for a single image

Image series annotation

Segmentation annotation with CellPose

"Human-in-the-loop" fine-tuning

- Run segmentation with CellPose Model
- Correct the segmentation manually
- Update the model (fine-tune on new annotations)
- (Repeat)

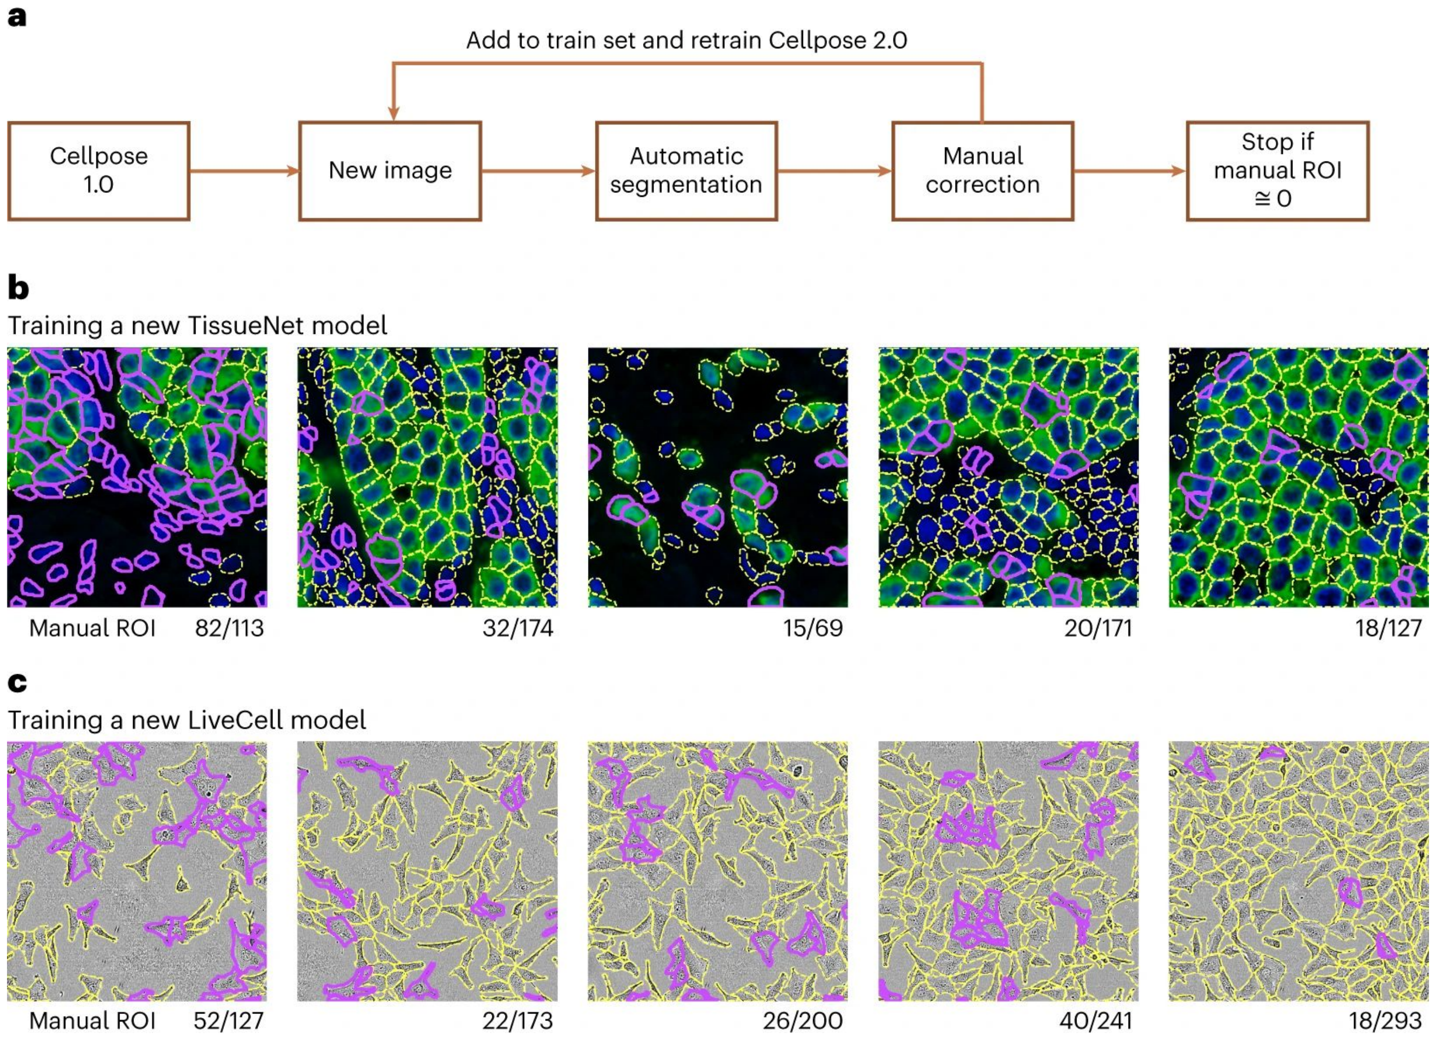

Segmentation annotation with CellPose

“Human-in-the-loop” fine-tuning

- Run segmentation with CellPose Model
- Correct the segmentation manually
- Update the model (fine-tune on new annotations)
- (Repeat)

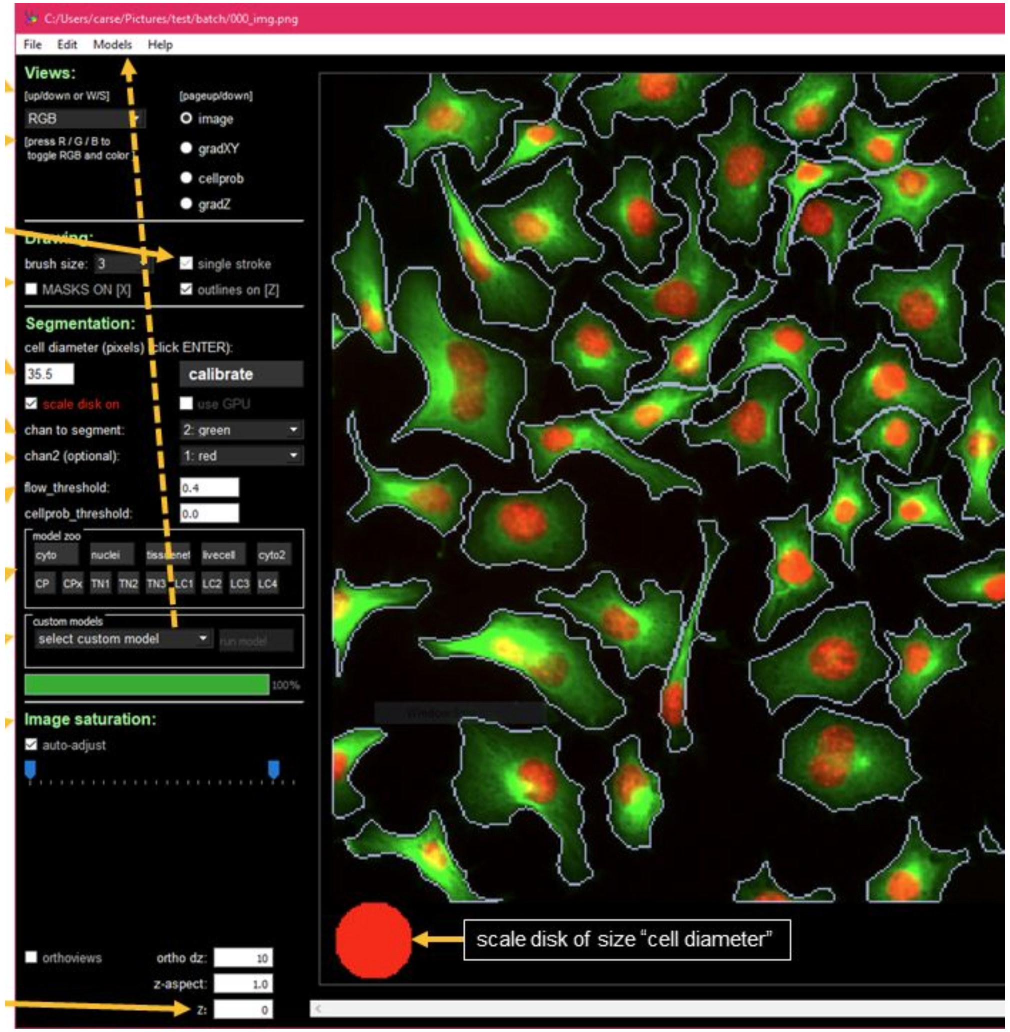

brush size: 3

☑ single stroke

☐ MASKS ON [X]

☑ outlines on [Z]

cell diameter (pixels) (click ENTER): 35.5

☑ scale disk on

☐ use GPU

chan to segment: 2: green

chan2 (optional): 1: red

flow_threshold: 0.4

cellprob_threshold: 0.0

☑ auto-adjust

☐ orthoviews

ortho dz: 10

z-aspect: 1.0

Z: 0

Classification annotation with µSAM

Live demo!

This is a prototype! (But can already be used)

Other tools for annotation

- QuPath:
  - For segmentation + classification
  - Very popular for histopathology

- Microscopy Image Browser
  - For segmentation
  - Very popular for electron microscopy

- Ilastik Object Classification Workflow for classification

- Custom napari tools for classification

Third Exercise & Project Preparation

Plan

- Exercise 3: Try different tools on the Covid-IF data:
  - StarDist Napari Plugin (Nucleus Segmentation)
  - µSAM (Napari Plugin)
  - CellPose
  - See here for details

- Projects / work on your own data:
  - Short (!) presentations in the afternoon.
  - We then discuss possible approaches and form groups.