GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN IN PUBLICA COMMODA SEIT 1737

COMPUTATIONAL
CELL ANALYTICS

Segment Anything for Microscopy

Interactive & Automatic Microscopy Segmentation

Constantin Pape

Institut für Informatik, Georg August Universität Göttingen

ccpape

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

Things the group does…

Method development for microscopy image analysis:
Vision: from images to insight and clinical relevance in collaboration with life scientists

Segmentation and tracking tasks

- High content microscopy for clinical decision making
- EM tomography for synaptic biology
- Volume EM for tissue and whole organism analysis

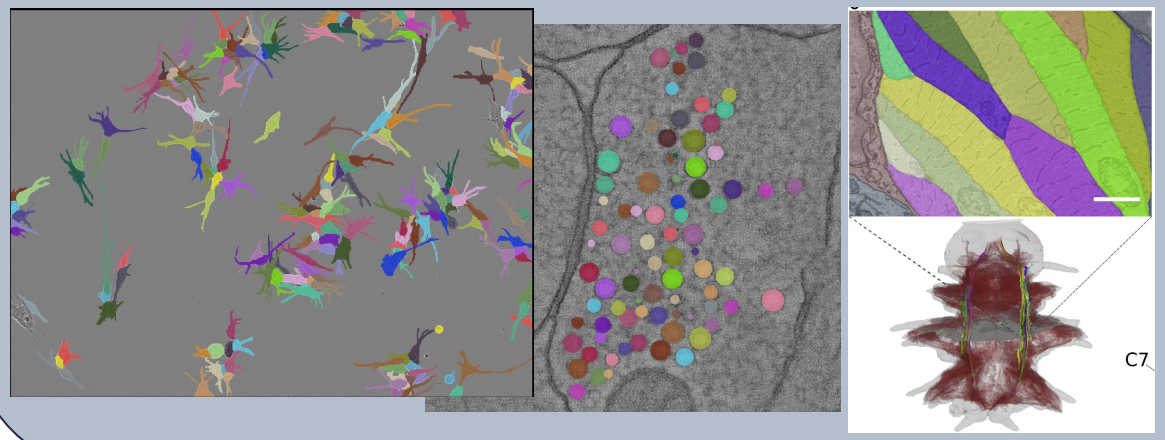

Representation learning for microscopy and multi-modal data

Protein structure analysis in cryo ET and **optical microscopy**

Segment Anything

Segment Anything

https://arxiv.org/abs/2304.02643

Pretrained model for interactive segmentation from Meta.AI

SAM: Interactive segmentation

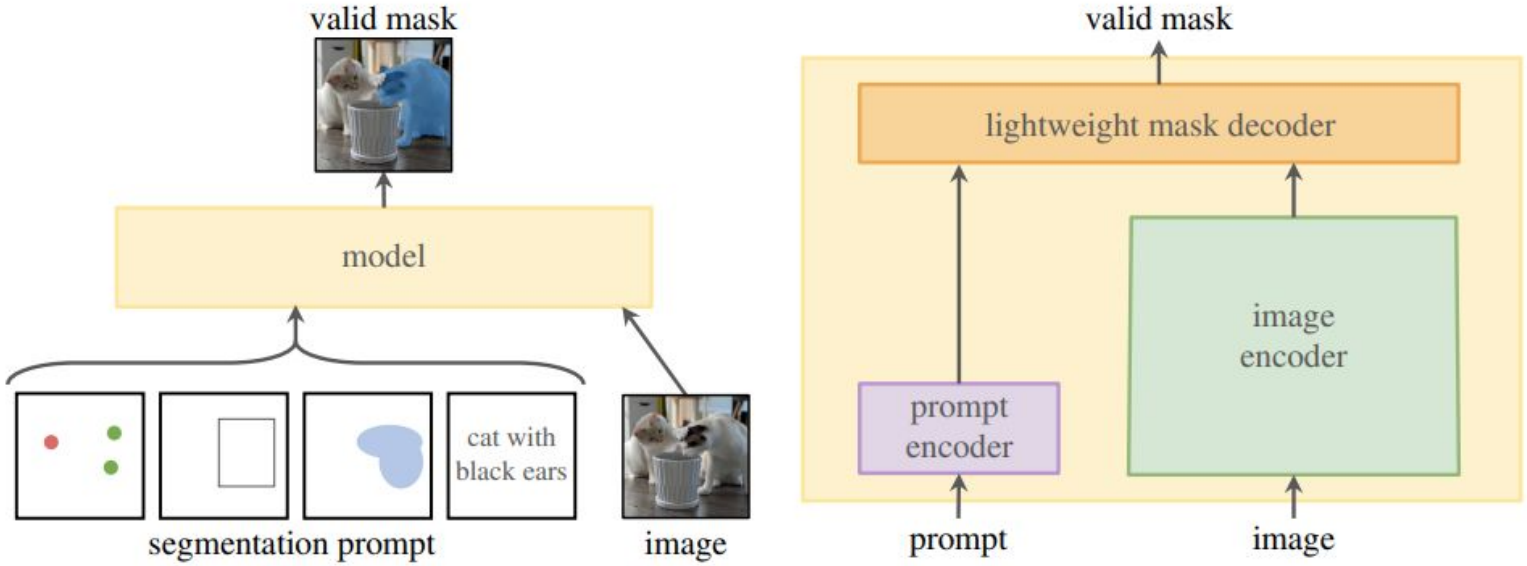

Segment Anything

https://arxiv.org/abs/2304.02643
* MobileSAM:
https://arxiv.org/abs/2306.14289

Pretrained model for interactive segmentation from Meta.AI

SAM: Interactive segmentation

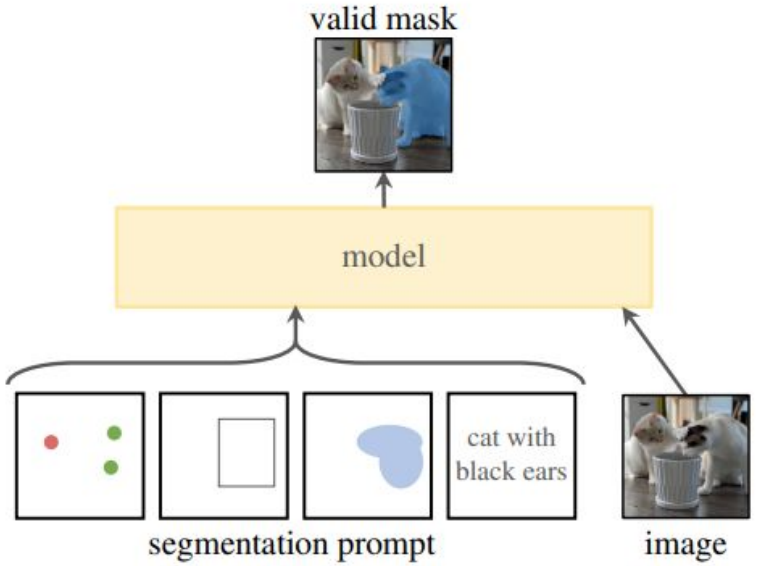

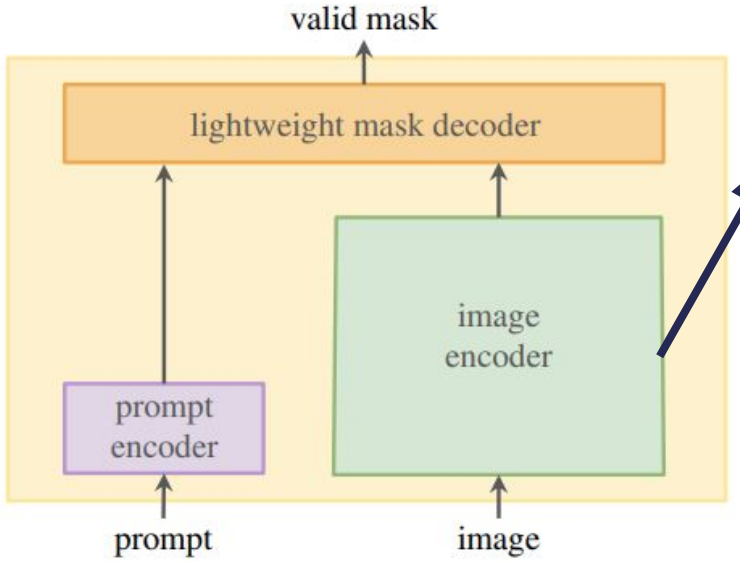

4 different sizes:
- VIT-B (Base)
- VIT-L (Large)
- VIT-H (Huge)
- VIT-T (Tiny)*

Segment Anything: What's special?

- Interactive segmentation: segment arbitrary objects from annotations
  - "prompts": points and/or box and/or mask
  - more prompts improve the predictions

- Versatile: can be integrated within pipelines that provide prompts
  - From user inputs, object detectors, nucleus seeds, …
  - Model is fully open-source!

Segment Anything: What's special?

- Interactive segmentation: segment arbitrary objects from annotations
  - "prompts": points and/or box and/or mask
  - more prompts improve the predictions

- Versatile: can be integrated within pipelines that provide prompts
  - From user inputs, object detectors, nucleus seeds, …
  - Model is fully open-source!

- **How?**
  - **Large dataset with diverse images and objects**
  - Iterative training loop

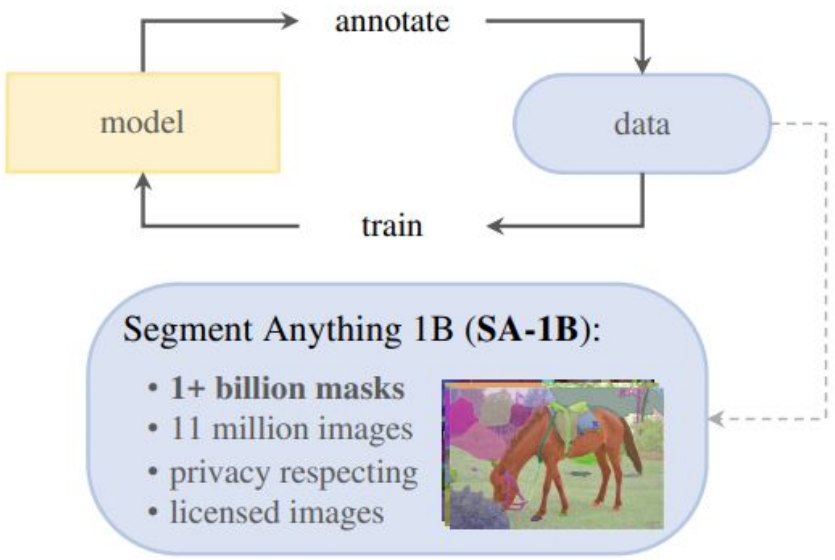

Segment Anything: Training iteration

Given image and ground-truth mask

- Compute image embeddings, sample positive point or box

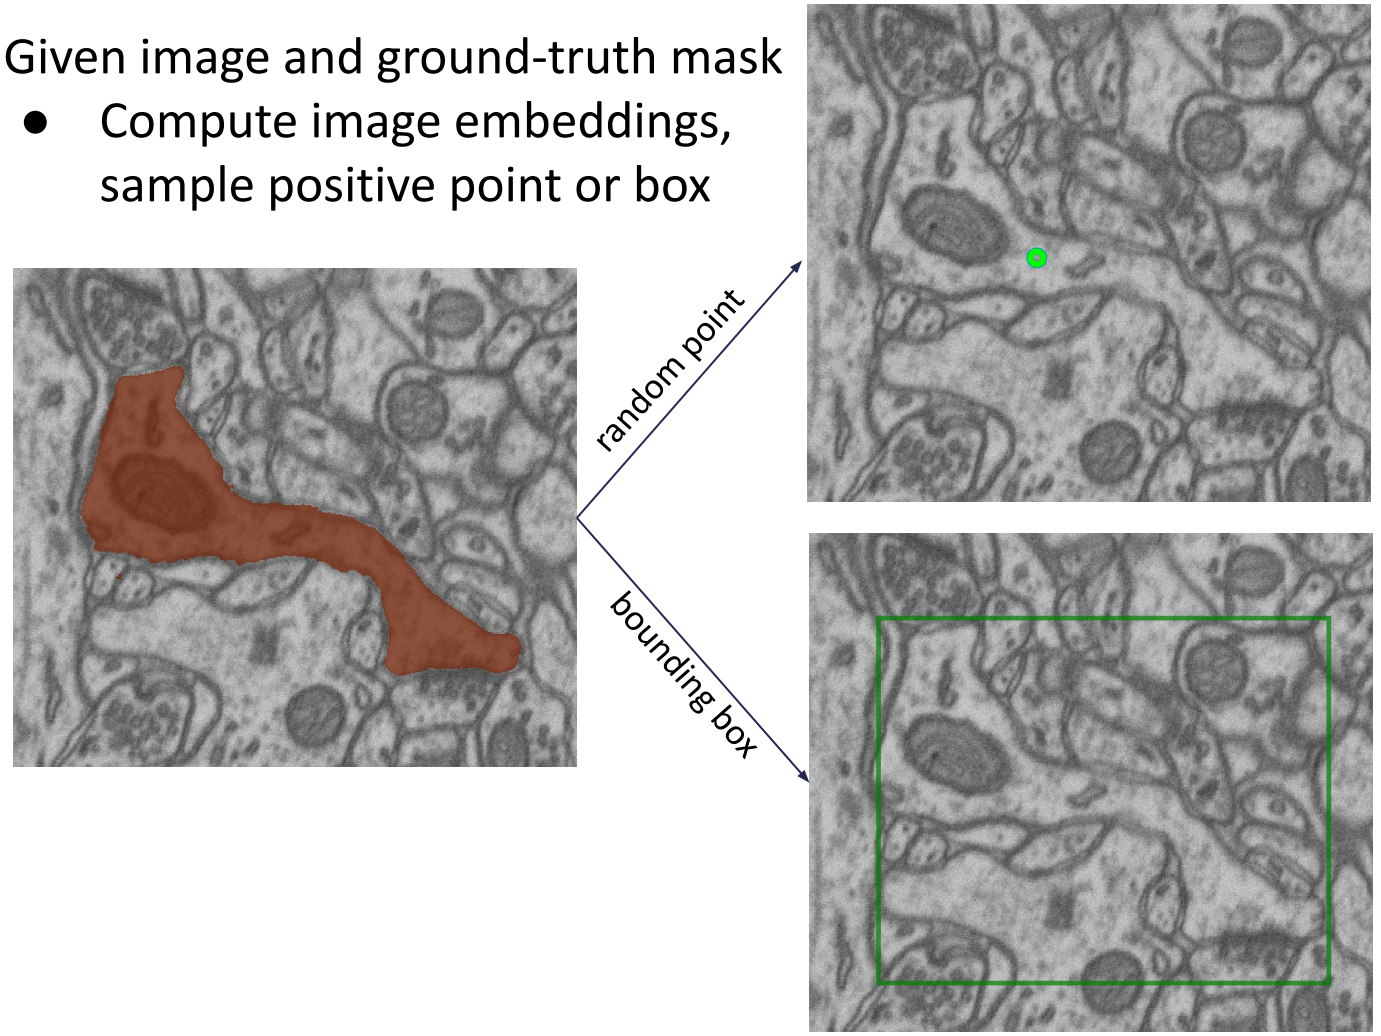

Segment Anything: Training iteration

Given image and ground-truth mask

- Compute image embeddings, sample positive point or **box**

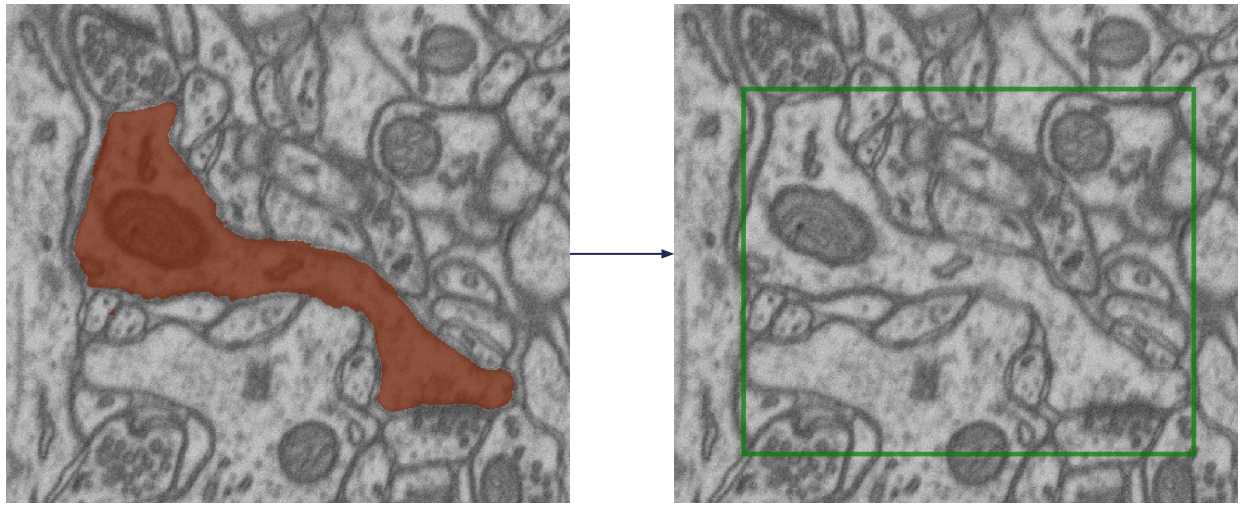

Segment Anything: Training iteration

Given image and ground-truth mask

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

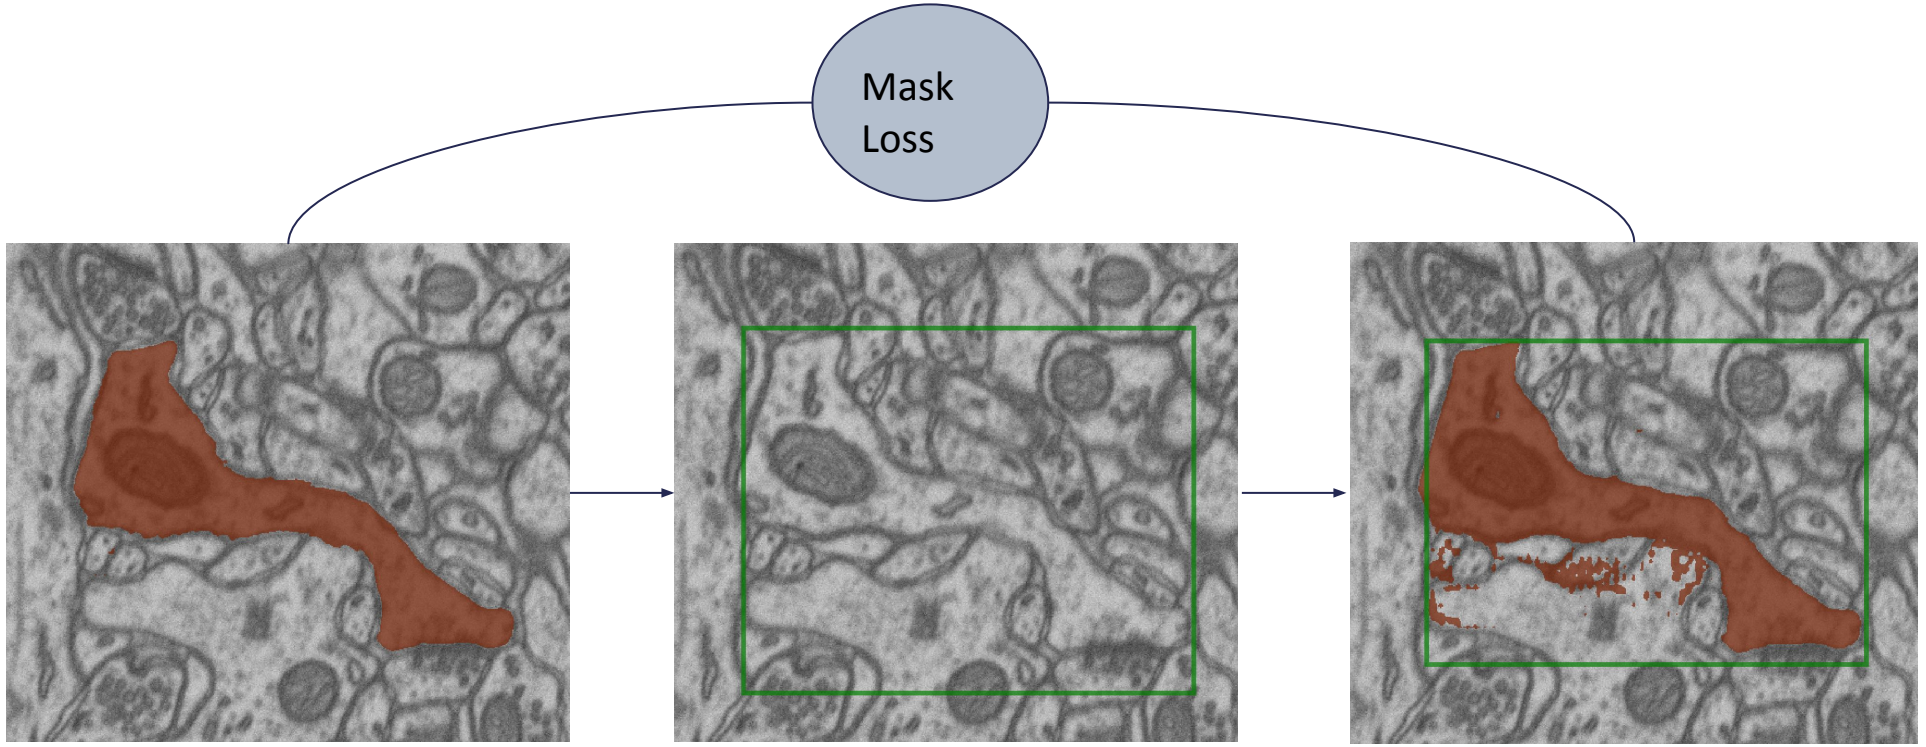

Segment Anything: Training iteration

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

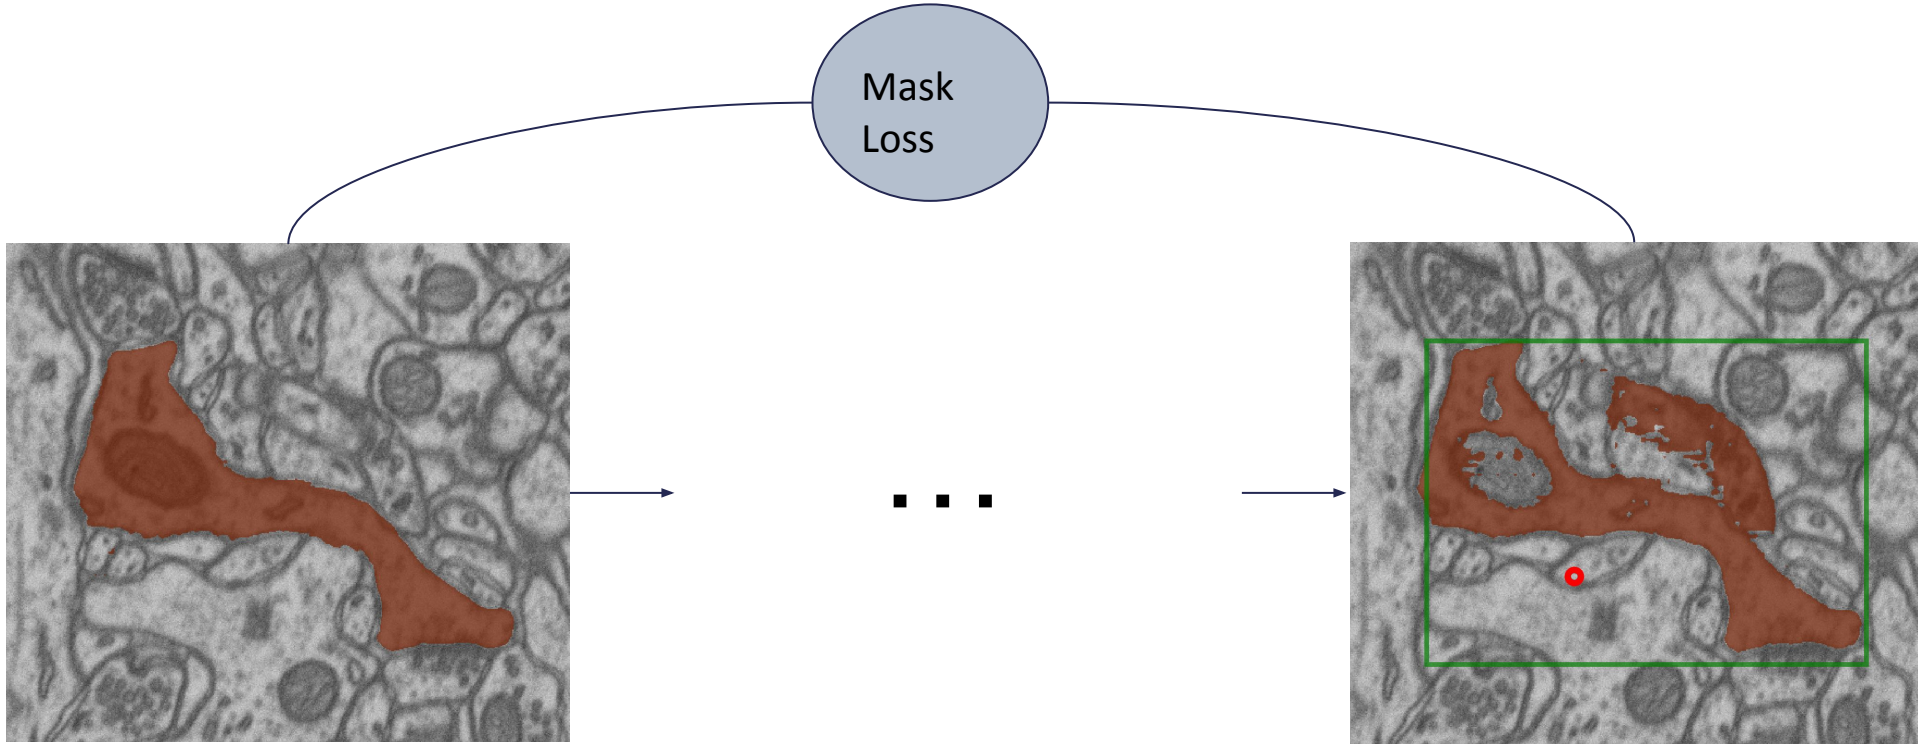

Segment Anything: Training iteration

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

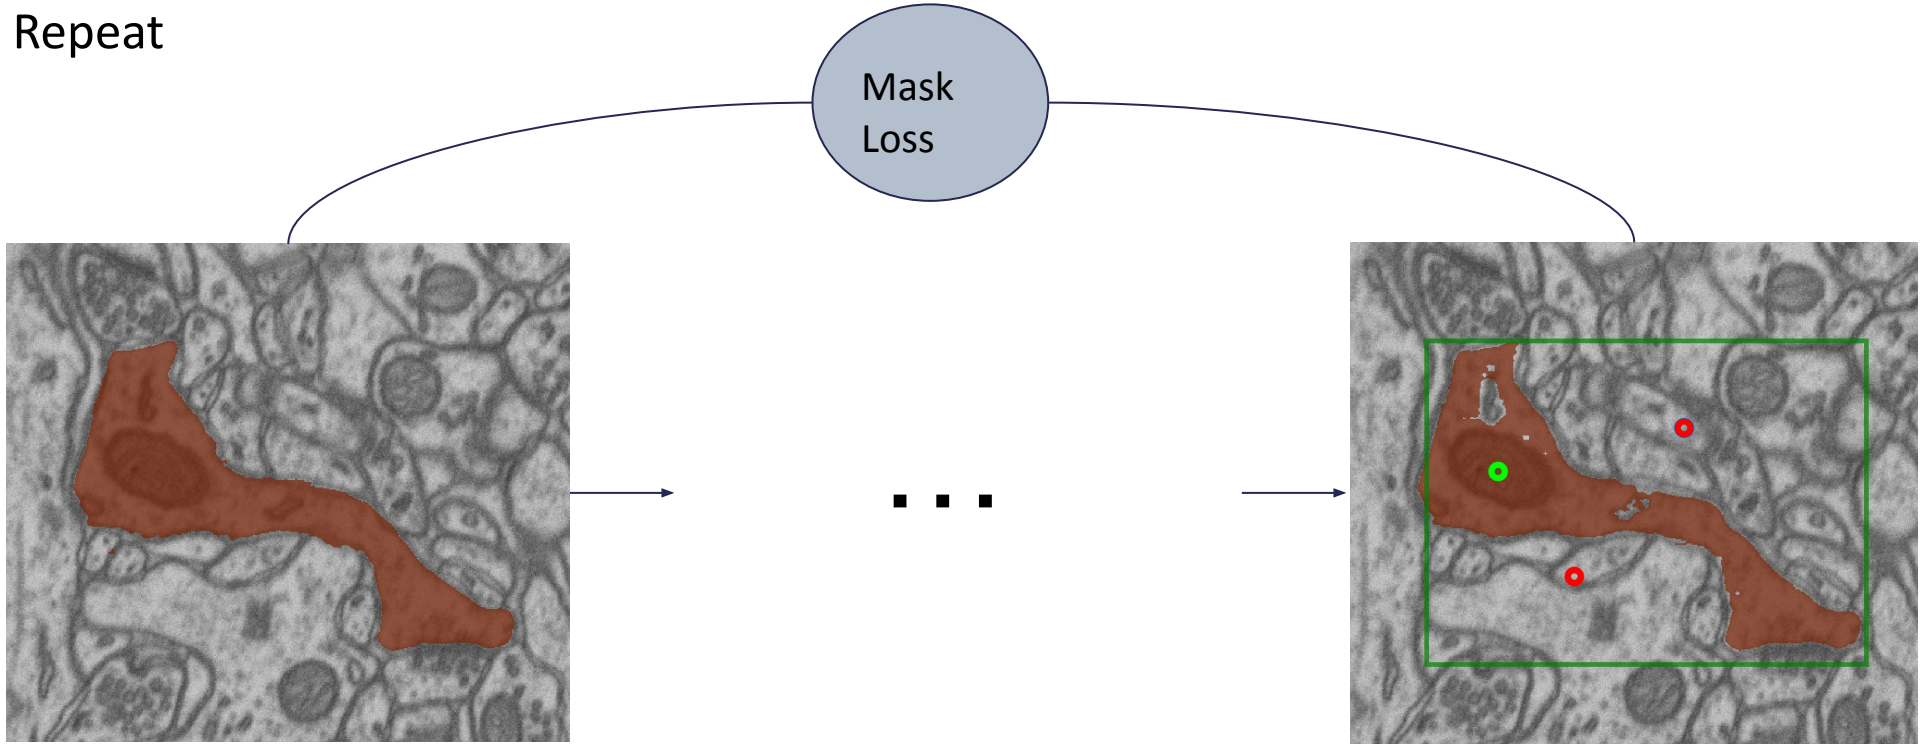

Segment Anything: Training iteration

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

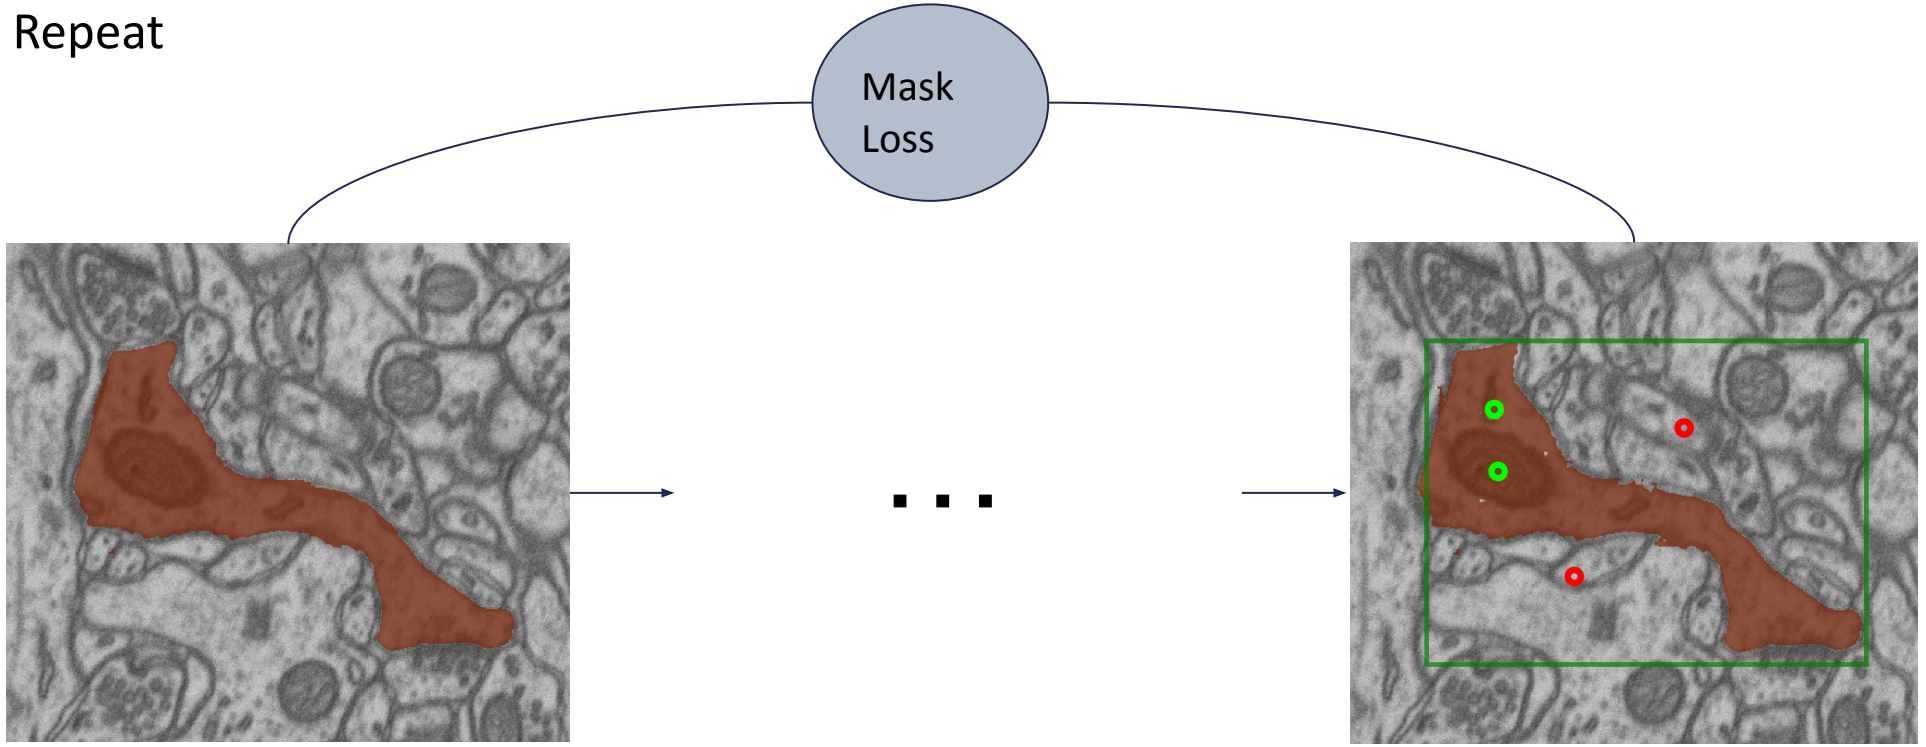

Segment Anything: Training iteration

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

Segment Anything: Capabilities

https://segment-anything.com/

Segmentation from user inputs (prompts)

Segment Anything: Capabilities

https://segment-anything.com/

Segmentation from user inputs (prompts)

Automatic Mask Generation (AMG)

Segment Anything for Microscopy

Anwai
Archit

Our aims & contributions

Archit, …, **Pape**, *bioRxiv (2023)*
https://doi.org/10.1101/2023.08.21.554208

- How well does SAM work for microscopy data? Which model size is best?
- Can we improve it (by finetuning) on microscopy data?
- Build a napari-based tool for interactive and automatic segmentation and tracking.

Collaboration between my group and DFKI; + several open source contributions.

Anwai
Archit

Our aims & contributions

Archit, …, **Pape**, *bioRxiv (2023)*
https://doi.org/10.1101/2023.08.21.554208

- How well does SAM work for microscopy data? Which model size is best?

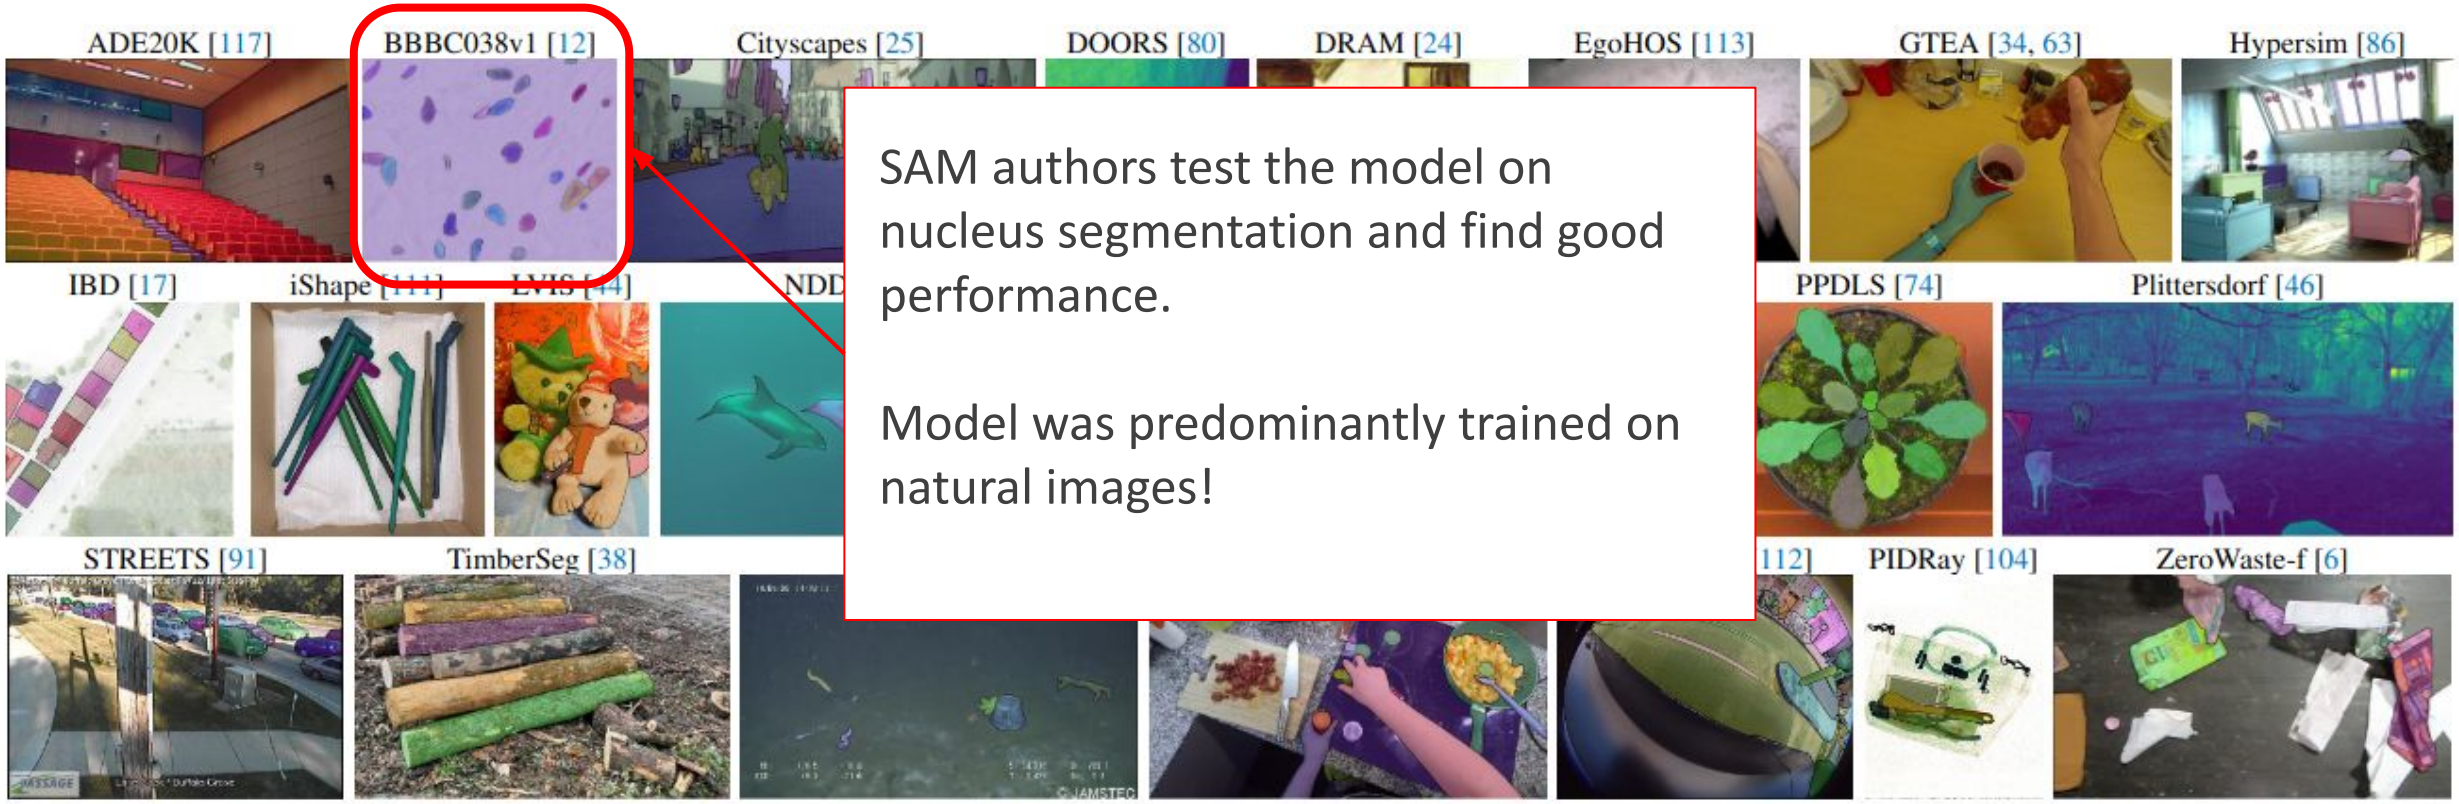

SAM authors test the model on nucleus segmentation and find good performance.

Model was predominantly trained on natural images!

Anwai
Archit

Our aims & contributions

Archit, …, **Pape**, *bioRxiv (2023)*
https://doi.org/10.1101/2023.08.21.554208

- How well does SAM work for microscopy data? Which model size is best?
- Can we improve it by finetuning on microscopy data?
- Build a napari-based tool for interactive and automatic segmentation and tracking.

We are in revision; will submit revised version this week!

Results are from revision experiments and not in preprint yet.

Finetuning SAM

Our contributions:

- Re-implement iterative training
  - Original code not published
  - Complex procedure
  - Use to finetune **SAM components**

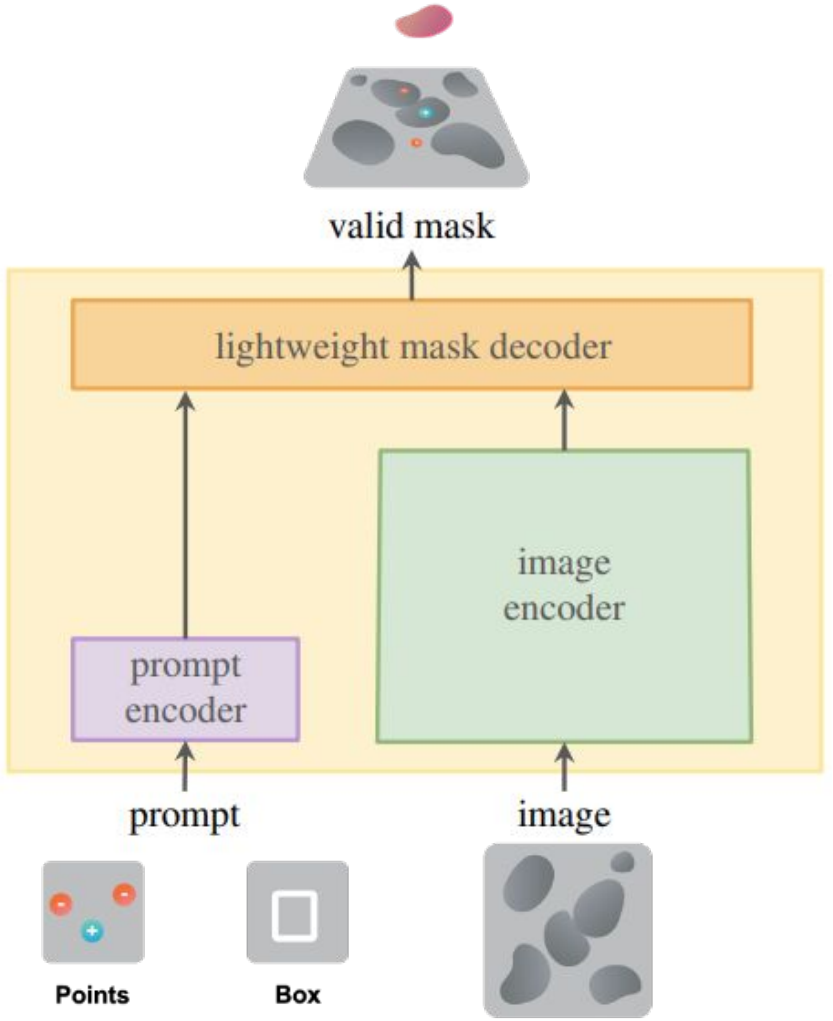

Finetuning SAM + improve instance seg

Our contributions:

- Re-implement iterative training
  - Original code not published
  - Complex procedure
  - Use to finetune SAM components

- Add decoder for instance segmentation (AIS)
  - Predicts foreground
  - Regresses distances to boundary + centroid
  - Input for watershed

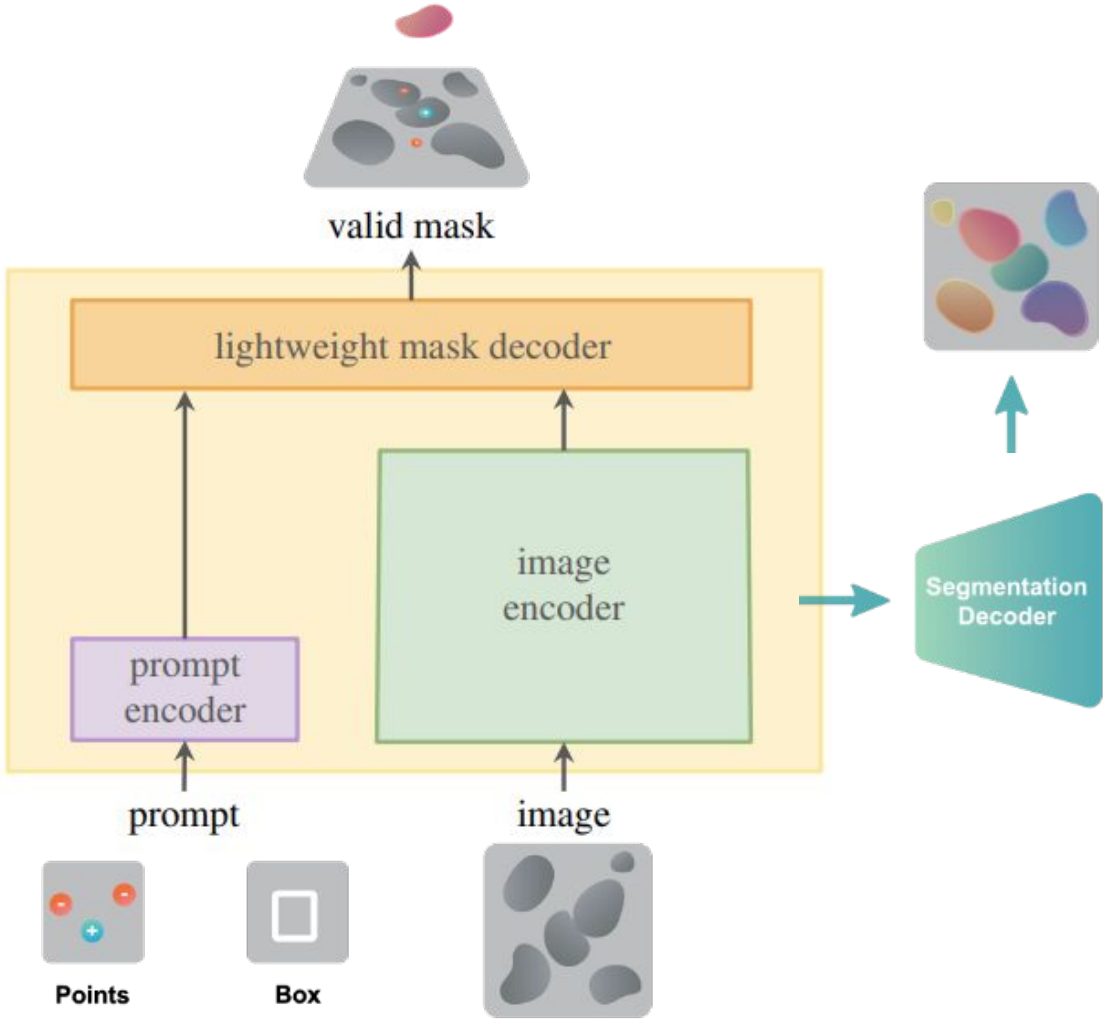

Finetuning for light microscopy

- Training data: cell and nucleus segmentation (published datasets)
  - Cells in Phase-contrast (LiveCELL)
  - Cells in Tissue (TissueNet)
  - Cells and Nuclei in Fluorescence (Neurips Cell Seg, DSB)
  - Cells in LightSheet (PlantSeg-Roots)
  - Bacteria in labelfree imaging (DeepBacs)

- Evaluate on test-split of training datasets (“in domain”)
  and unseen datasets (“out of domain”):
  - Nuclei and cells in confocal, cells in immunofluorescence, nuclei in histopathology, …

- Compare interactive and automatic instance segmentation
  - Compare to CellPose baseline for automatic segmentation

Interactive Segmentation: In domain & Out-of-domain

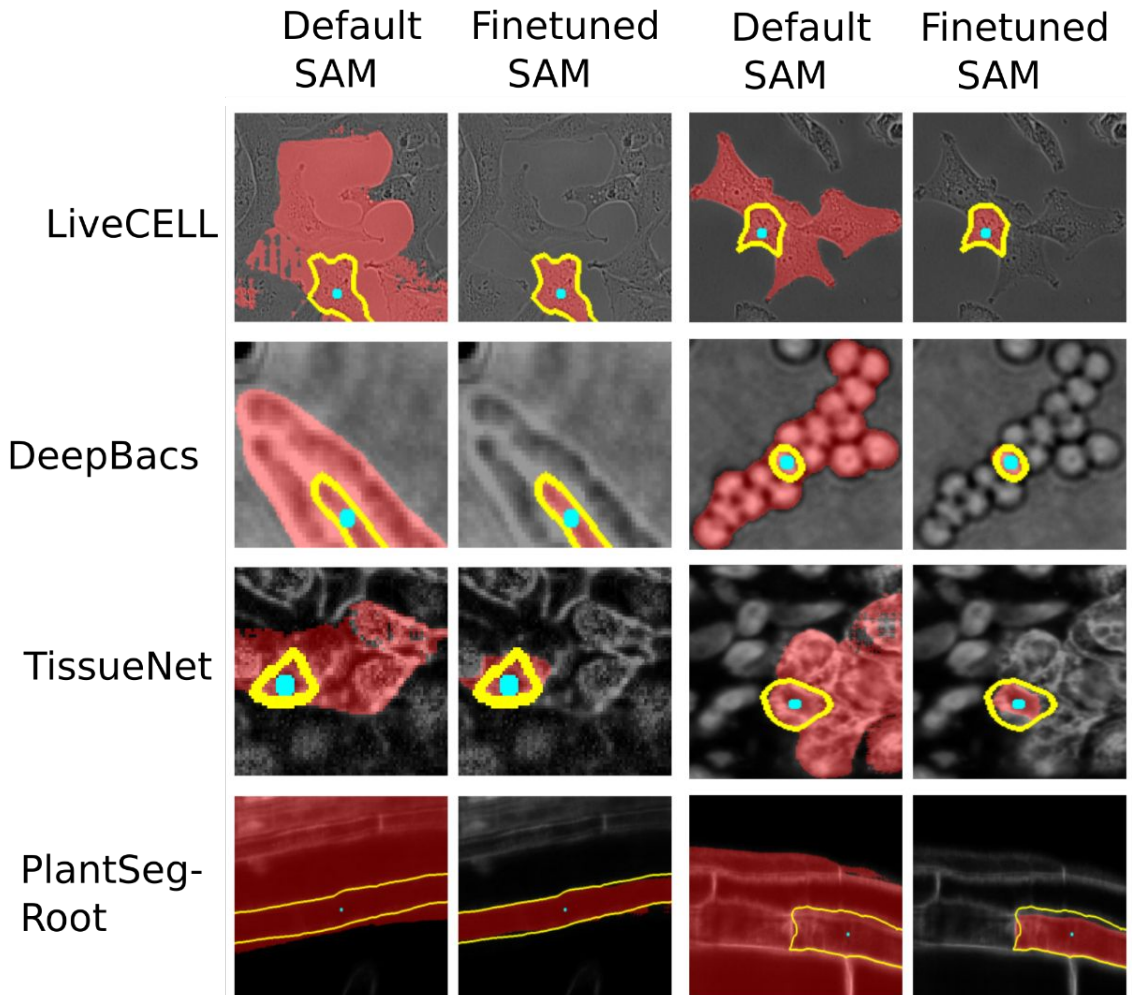

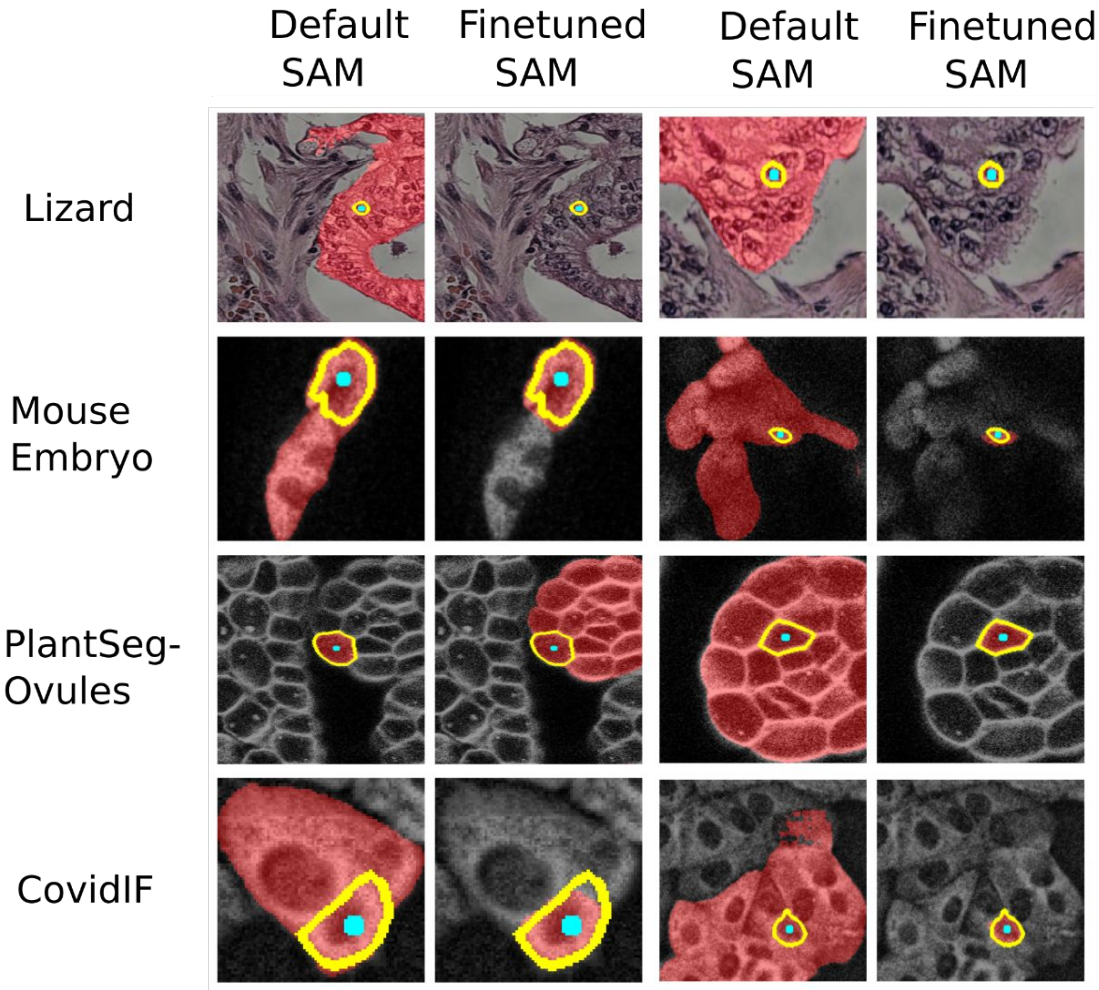

Results: In Domain

Results for LIVECell Dataset
(In Domain; Test Split)

Evaluation:

- Interactive Segmentation:
  - Derive prompts from ground-truth, improve iteratively
- Instance segmentation:
  - Compare with CellPose
- Both: compute segmentation accuracy (compared to ground-truth)

ViT Base

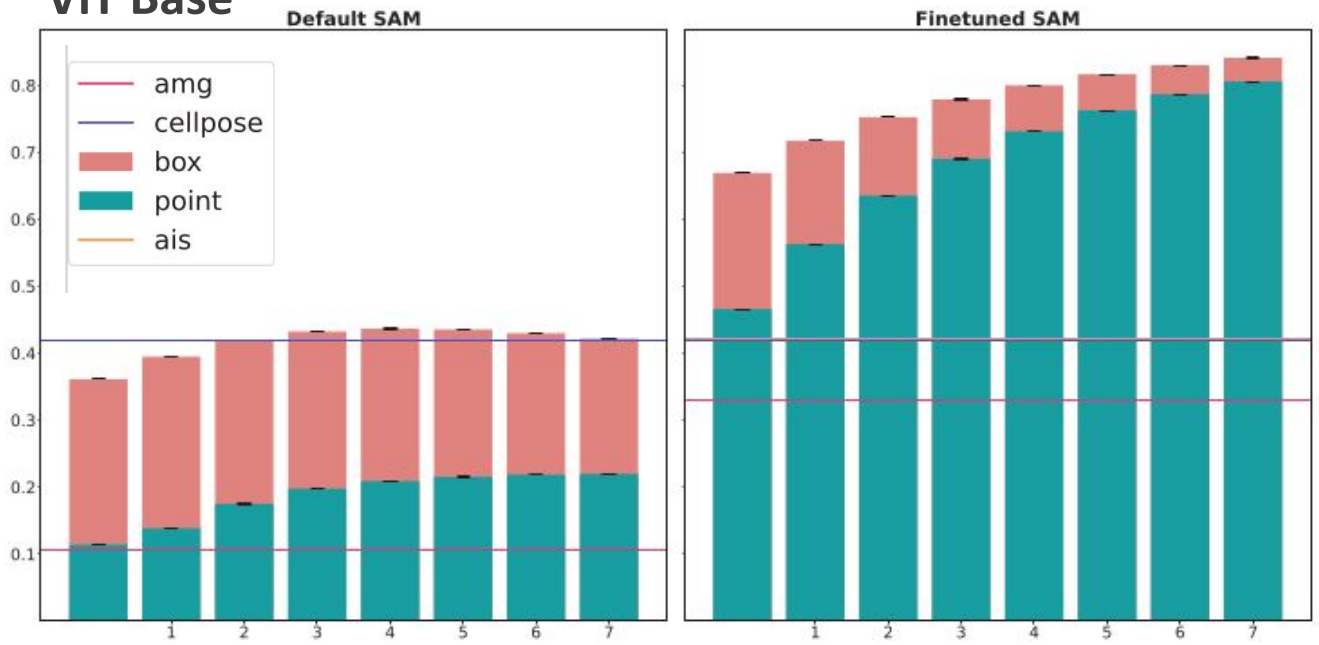

ViT Large

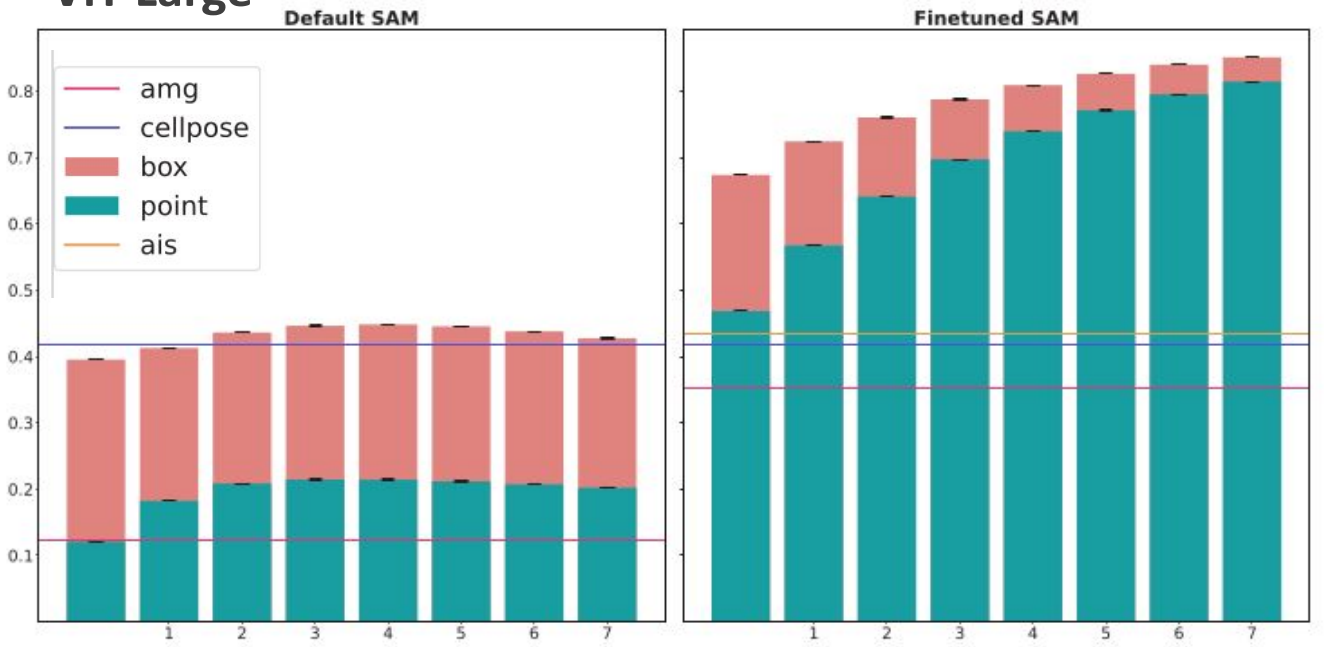

Automatic Segmentation

Instance segmentation on LIVECell Dataset

Runtimes on laptop (CPU);
including embedding computation (dominates for AIS)

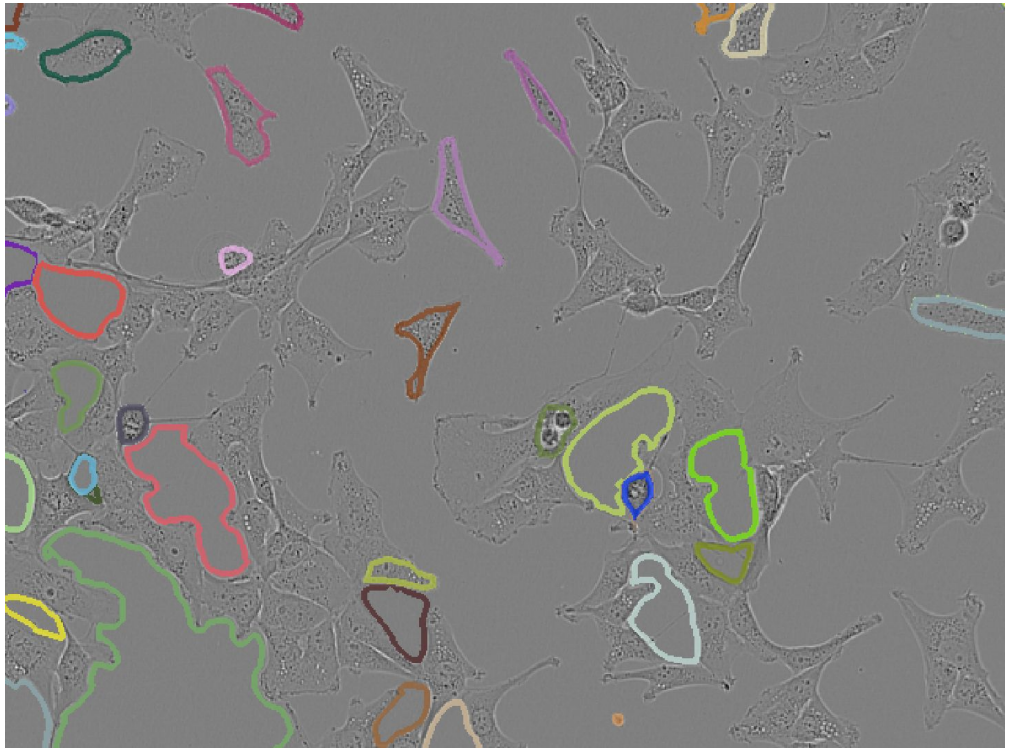

VIT-B
AMG: 75 sec

VIT-B-LM
AIS: 9 sec

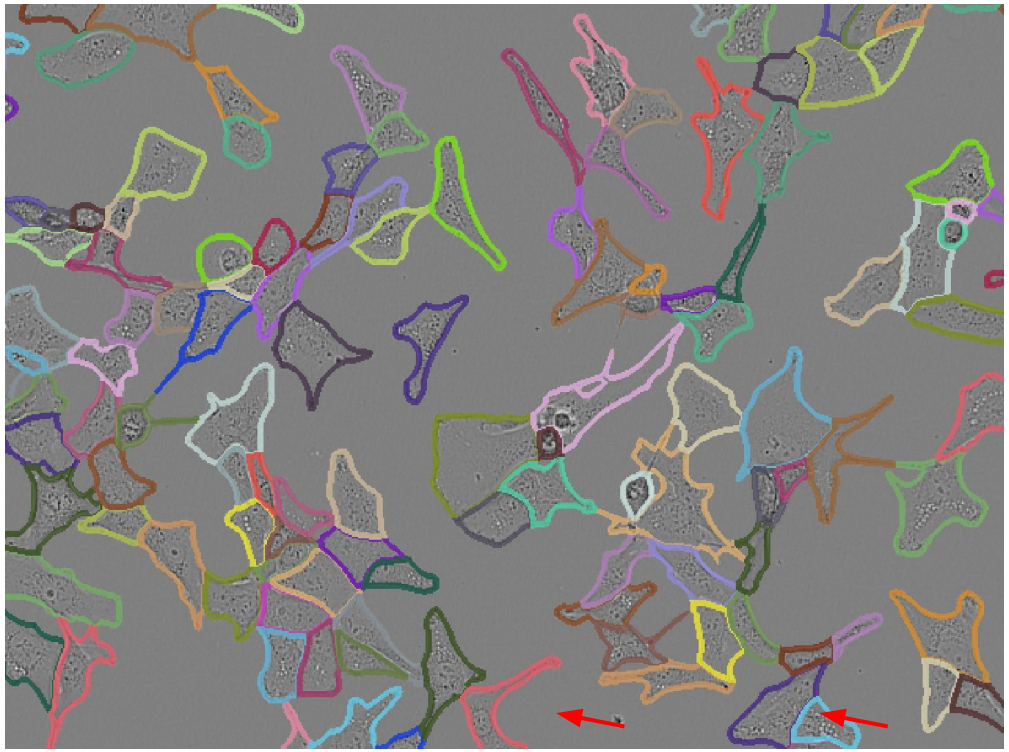

Results: Out of domain

Results for out of-domain datasets.

Same evaluation procedure as before.

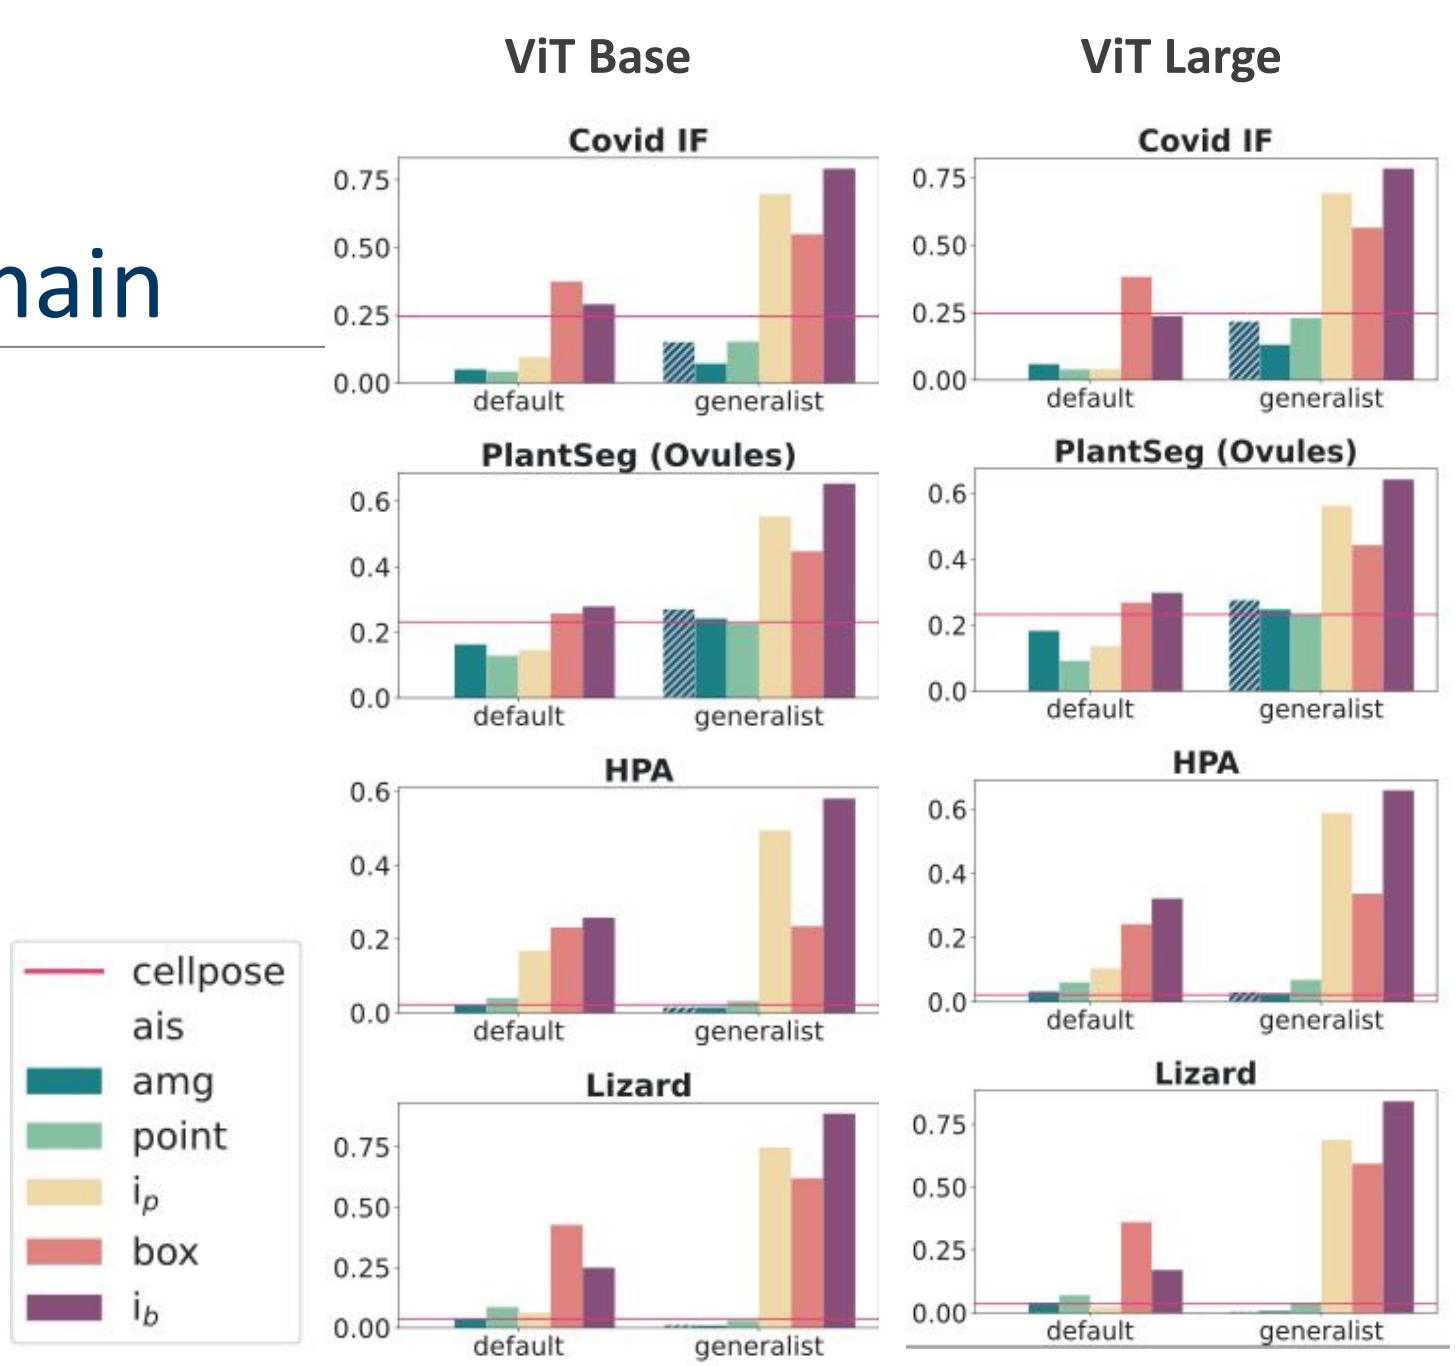

Results: Out of domain

Results for out of-domain datasets.

Same evaluation procedure as before.

Conclusions:

- Finetuning improves models!
- Best model: vit_l
  - If runtime matters: vit_b / vit_t
- Comparison to CellPose (automatic seg.):
  - Similar performance on most out of domain datasets (cyto2 model)

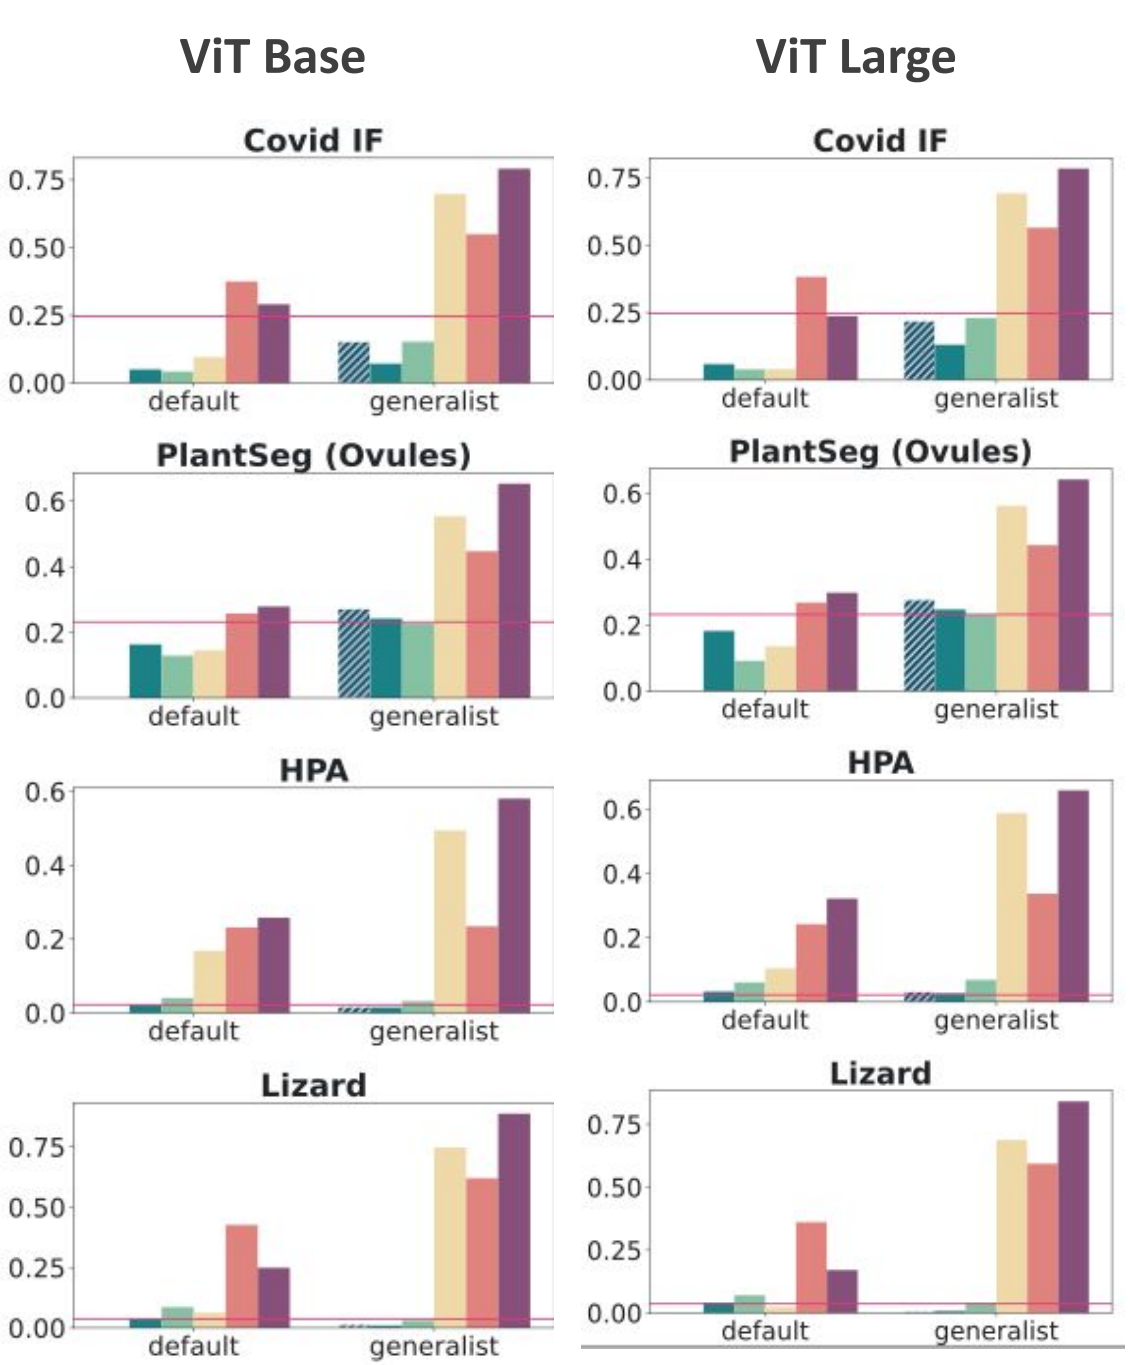

Finetuning for electron microscopy

- Training data: Mitochondria and nucleus segmentation in electron microscopy
  - Most training data from MitoNet (https://doi.org/10.1016/j.cels.2022.12.006).
- Compare default and finetuned model.
  - Compare automated segmentation with MitoNet.
- Evaluate on test-split of training datasets (“in domain”) and unseen datasets (“out of domain”)
  - Application to EM mitochondria from non-training data.

Interactive Segmentation: In domain & Out-of-domain

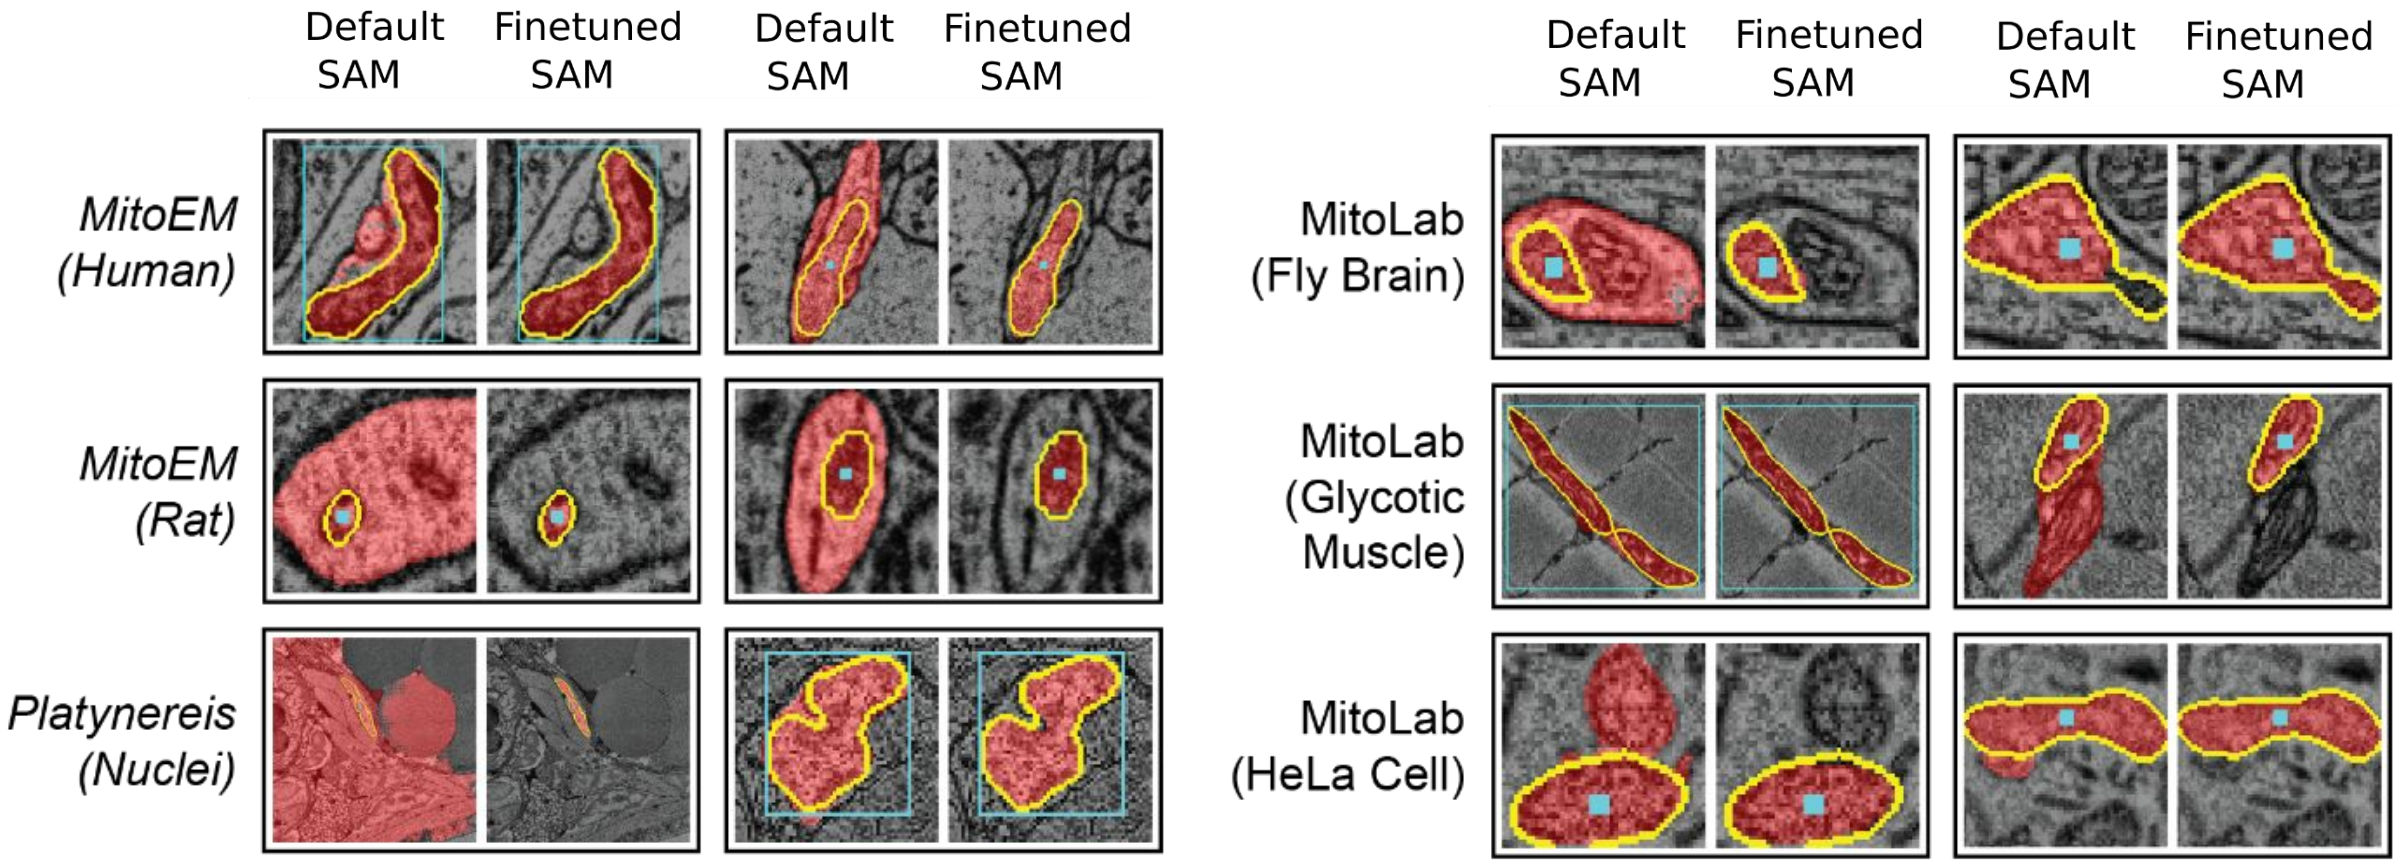

Results: In & out-of domain

Evaluation: Same approach as for LM

- In domain (top row)
- Out of domain (rest)

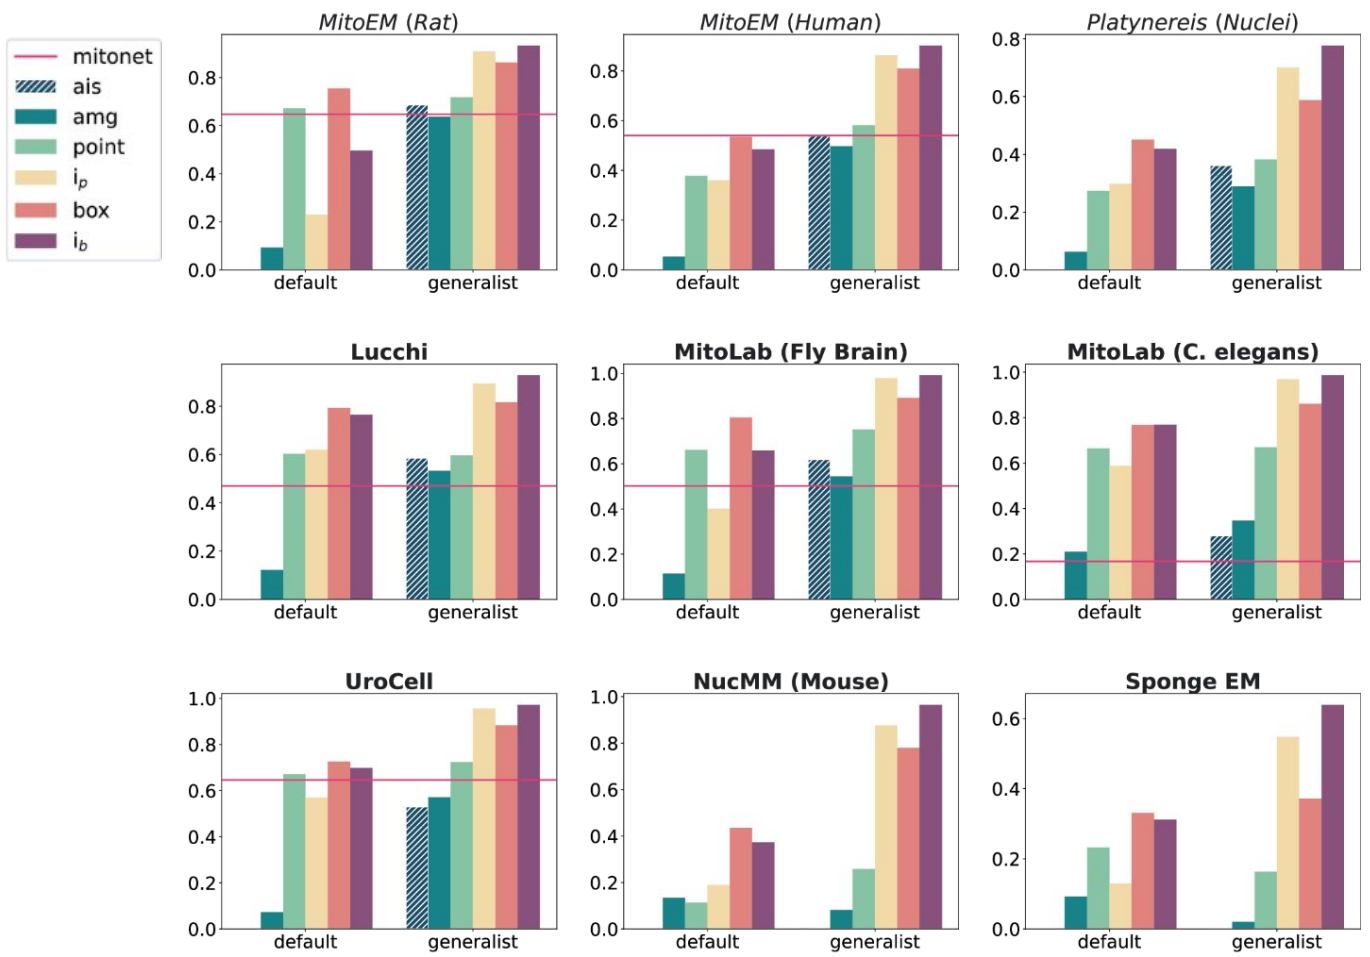

Results: In & out-of domain

Evaluation: Same approach as for LM

- In domain (top row)
- Out of domain (rest)

Conclusions:

- Finetuning improves, best model is vit_l
- Similar performance to MitoNet on most datasets (AIS)

- **Improves segmentation for some other organelles (cilia, microvilli), but worsens it for cellular compartments**
  - **Bigger diversity in EM!**

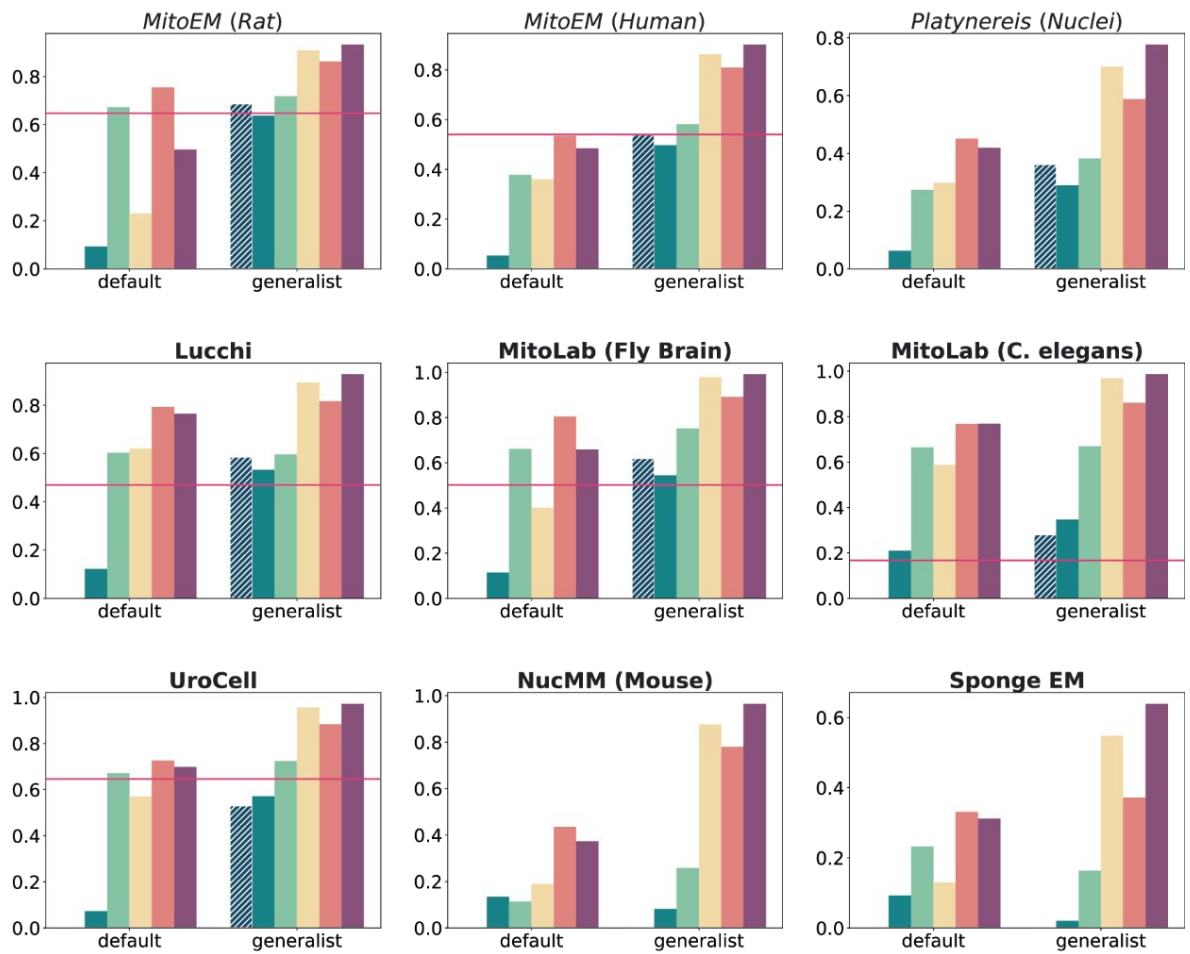

Mitochondria

Instance segmentation on Lucchi Dataset

Runtimes on laptop (CPU);
including embedding computation (dominates for AIS)

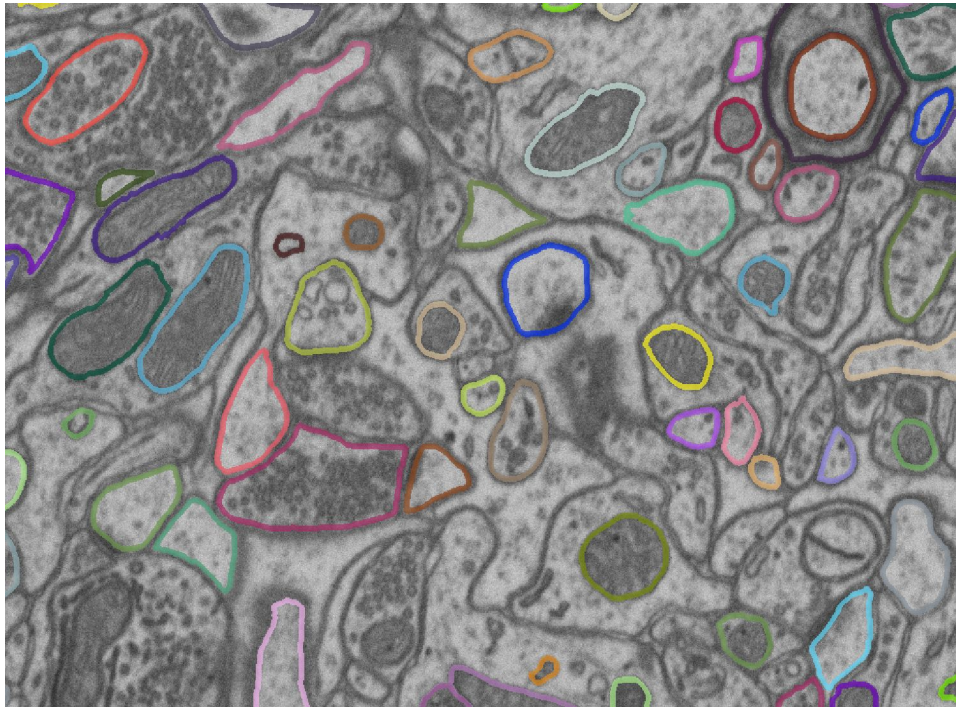

VIT-B
AMG: 80 sec

VIT-B-EM
AIS: 10 sec

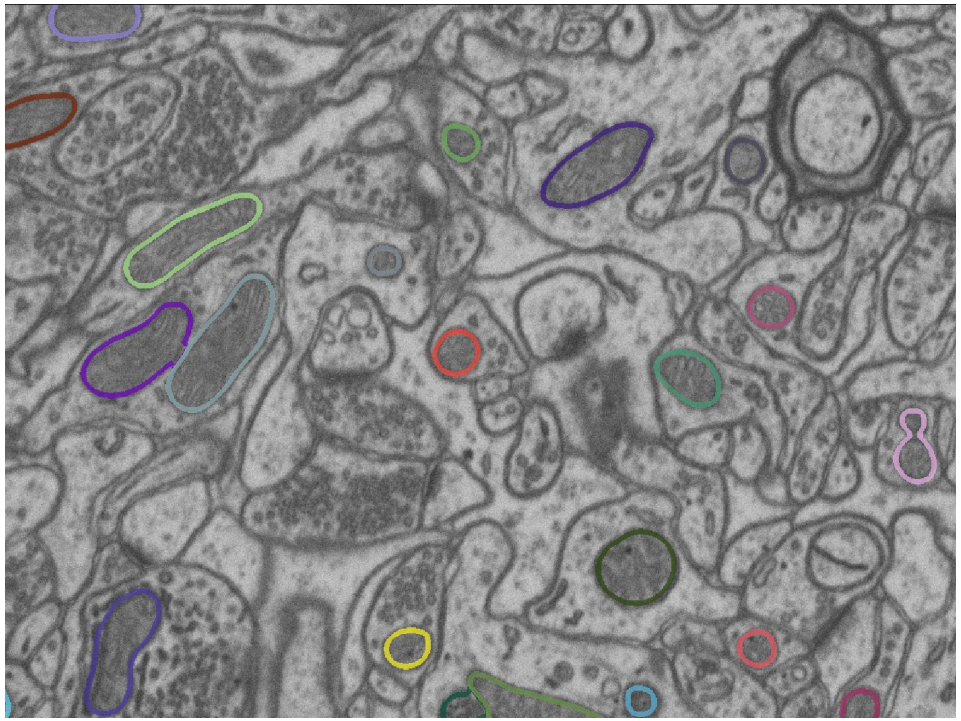

Finetuning as a user

CovidIF

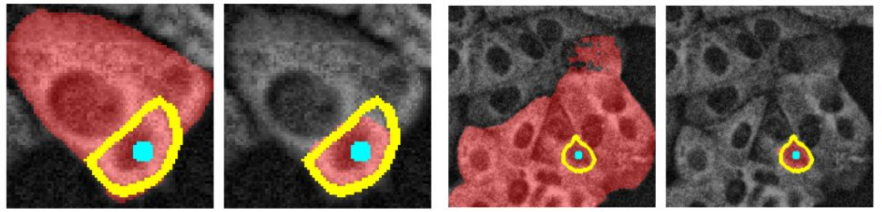

Improve models further for your data?

- How much data is needed?
- Which computational resources are required?

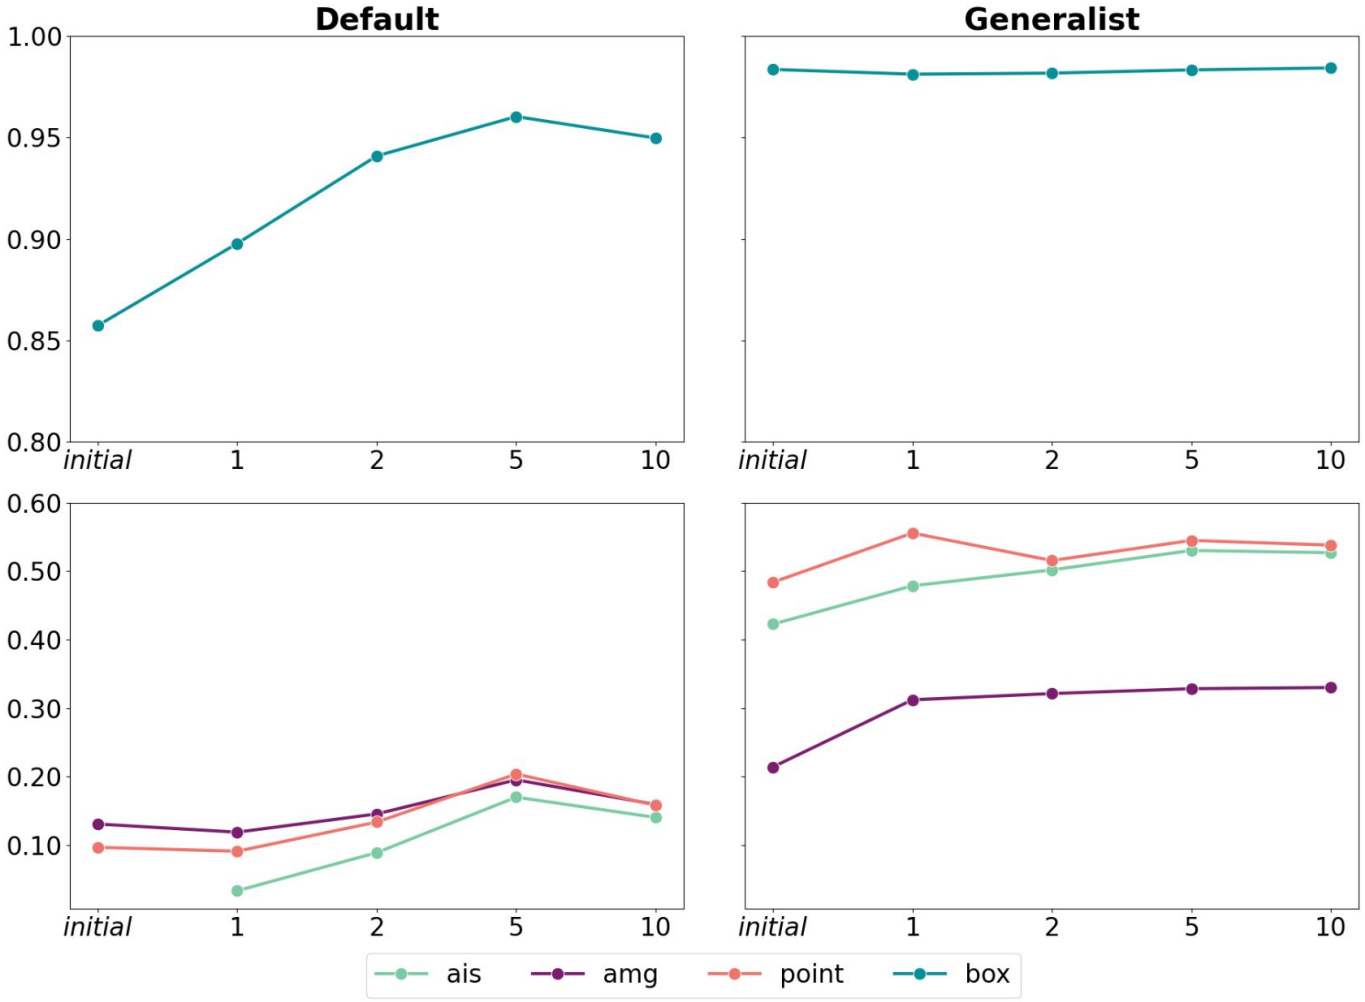

Finetuning as a user

CovidIF

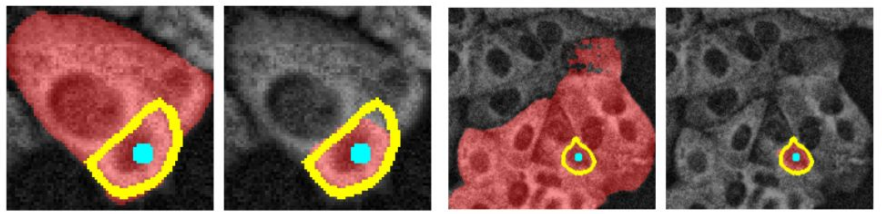

Improve models further for your data?

- How much data is needed?
- Which computational resources are required?

- Few images with annotations are sufficient!
- Finetuning is possible on CPU (but takes quite long); reasonable time on a GPU.

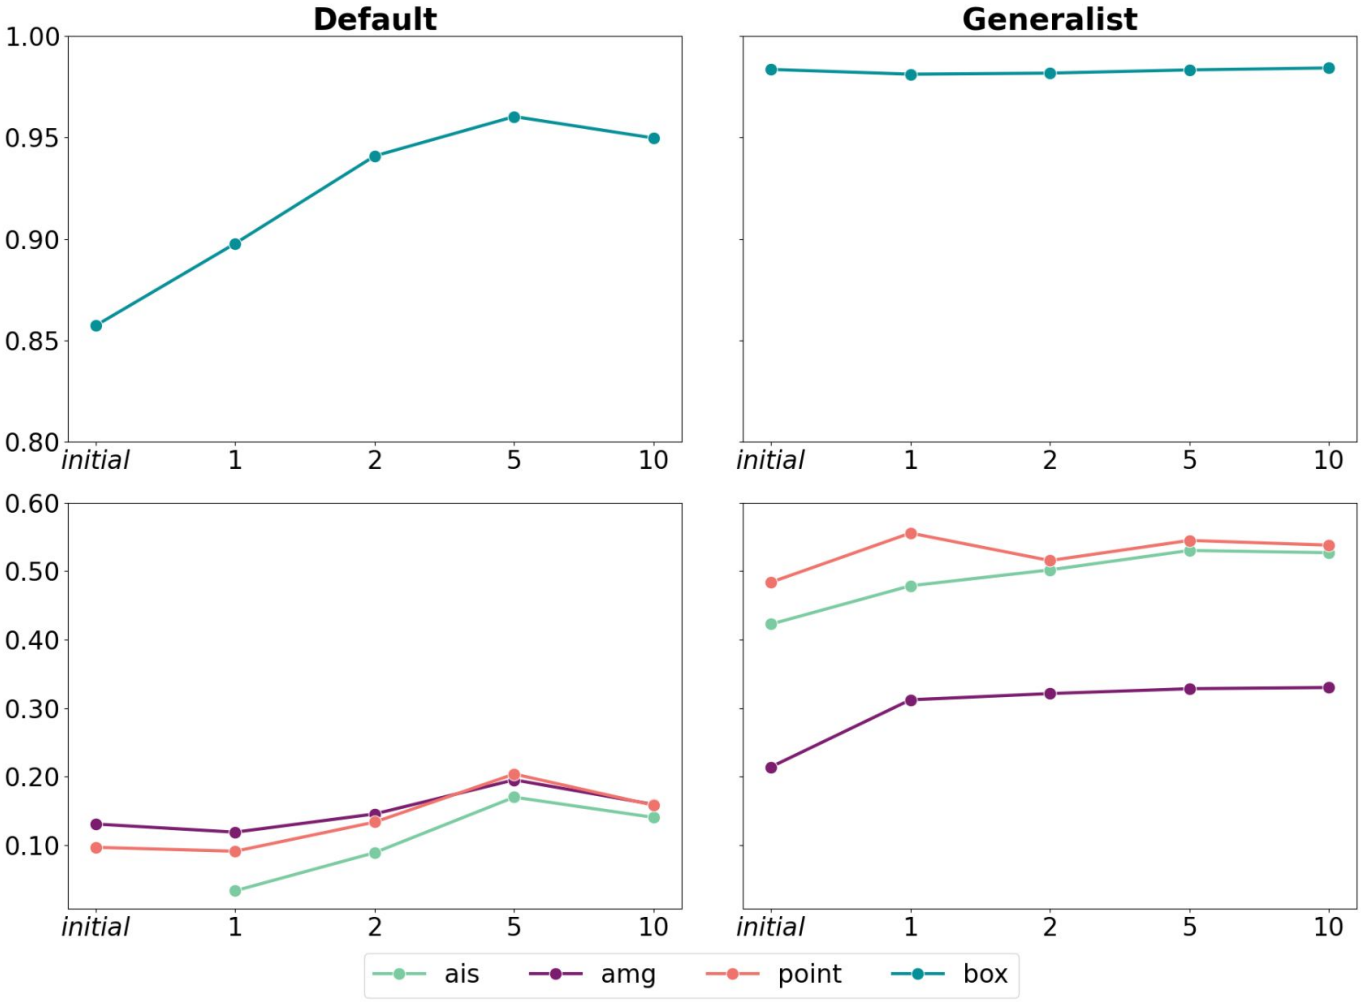

Application in practice

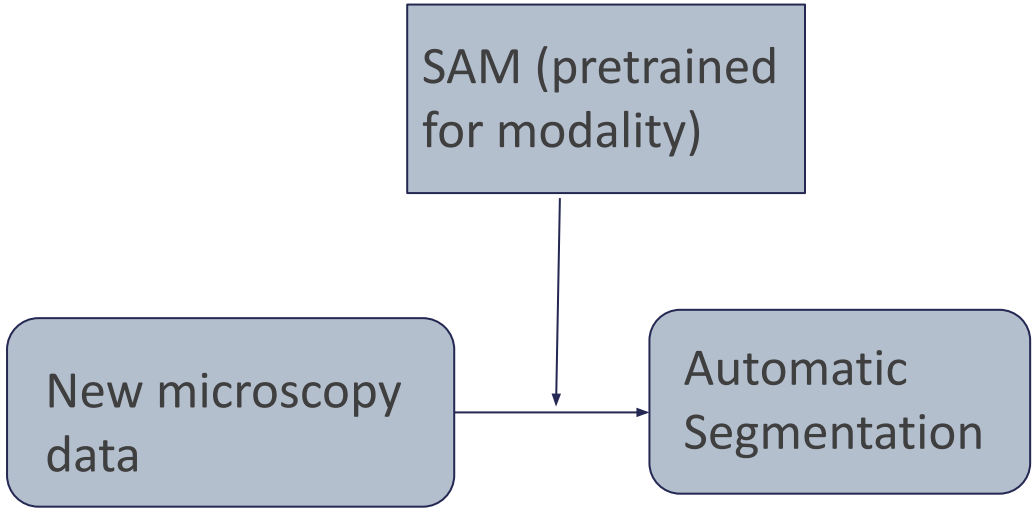

Application in practice

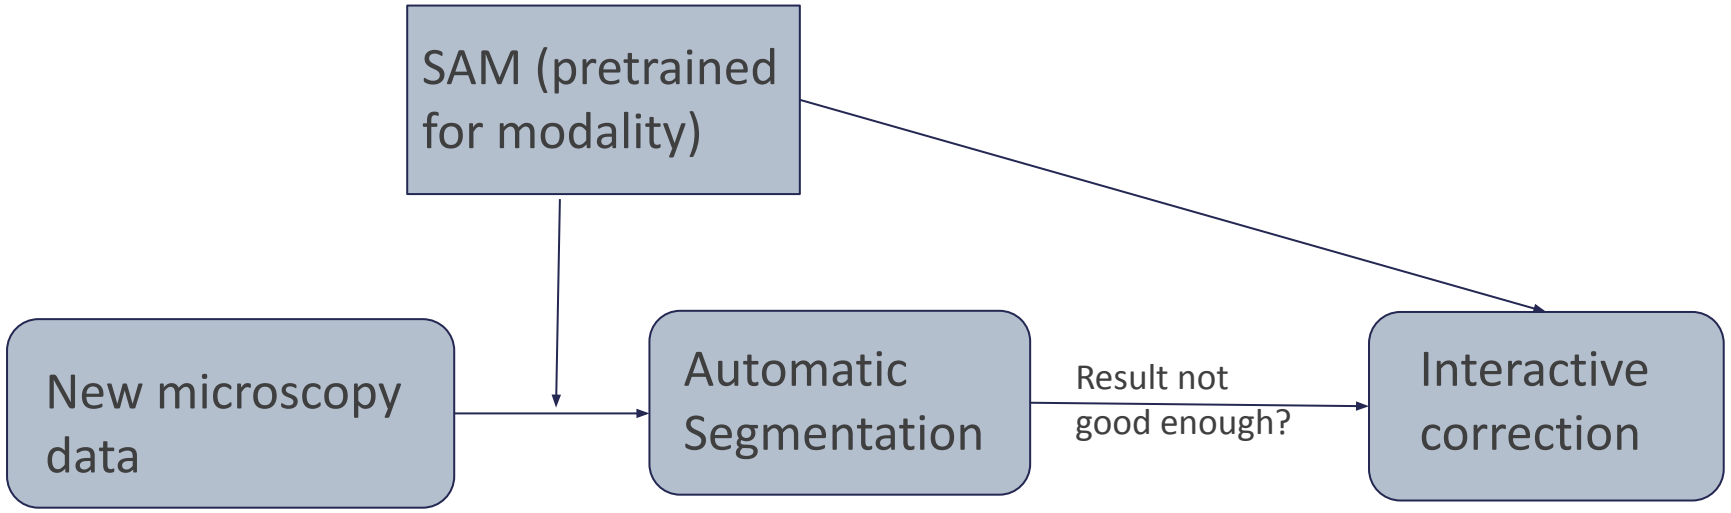

Application in practice

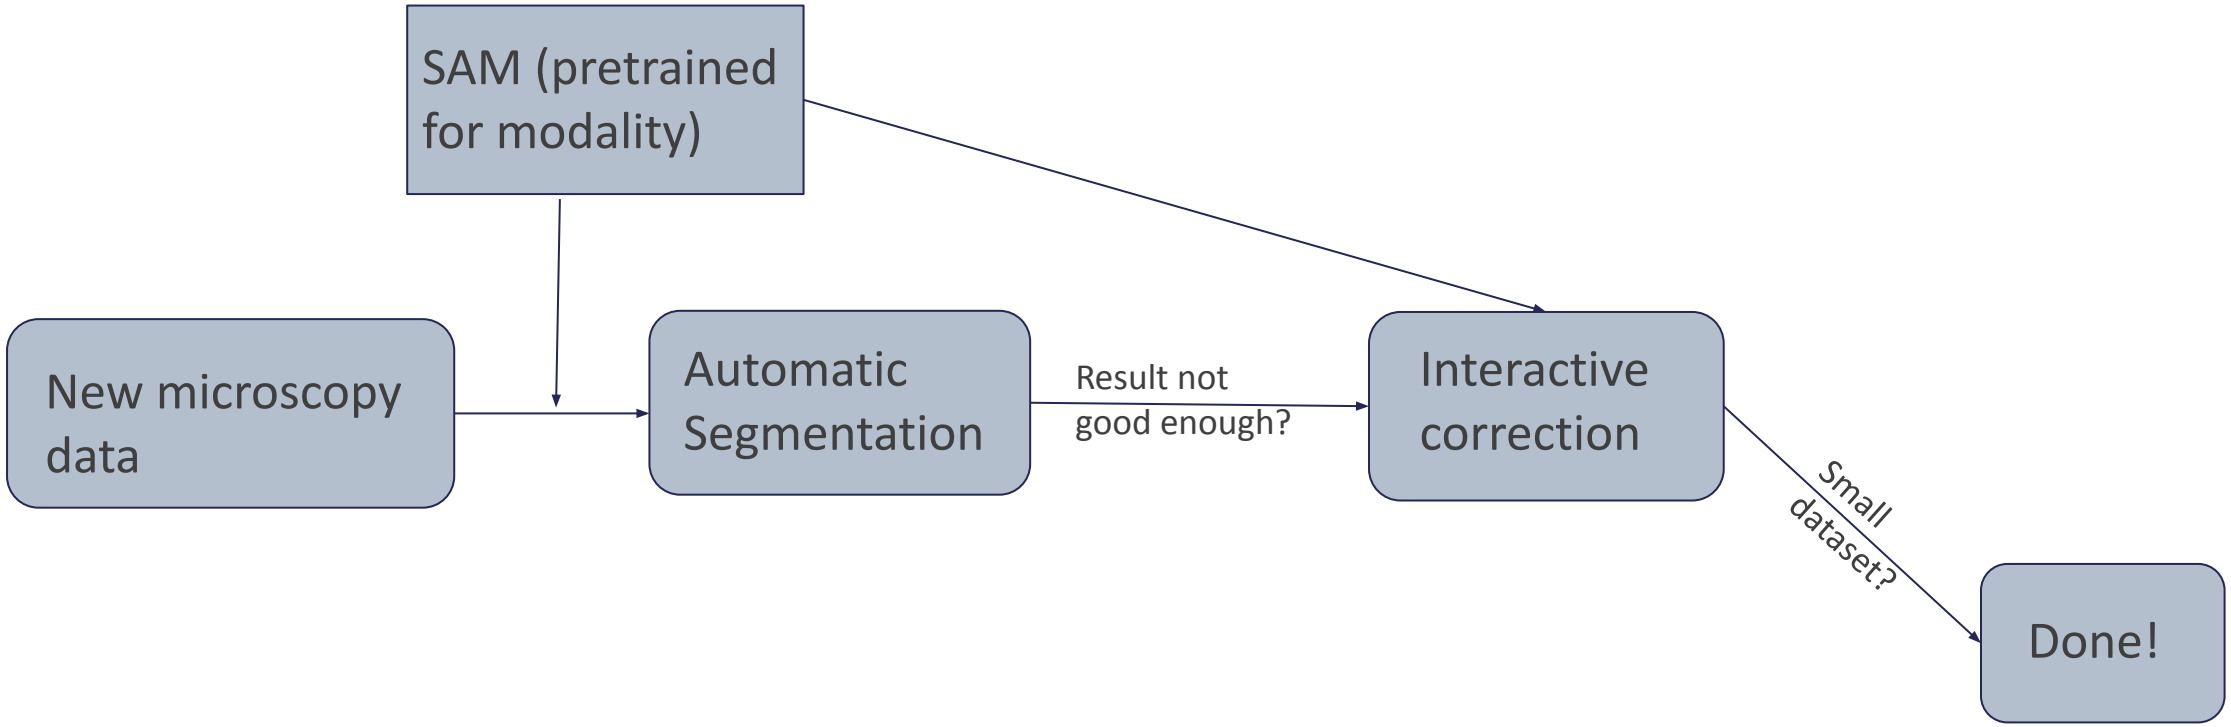

Application in practice

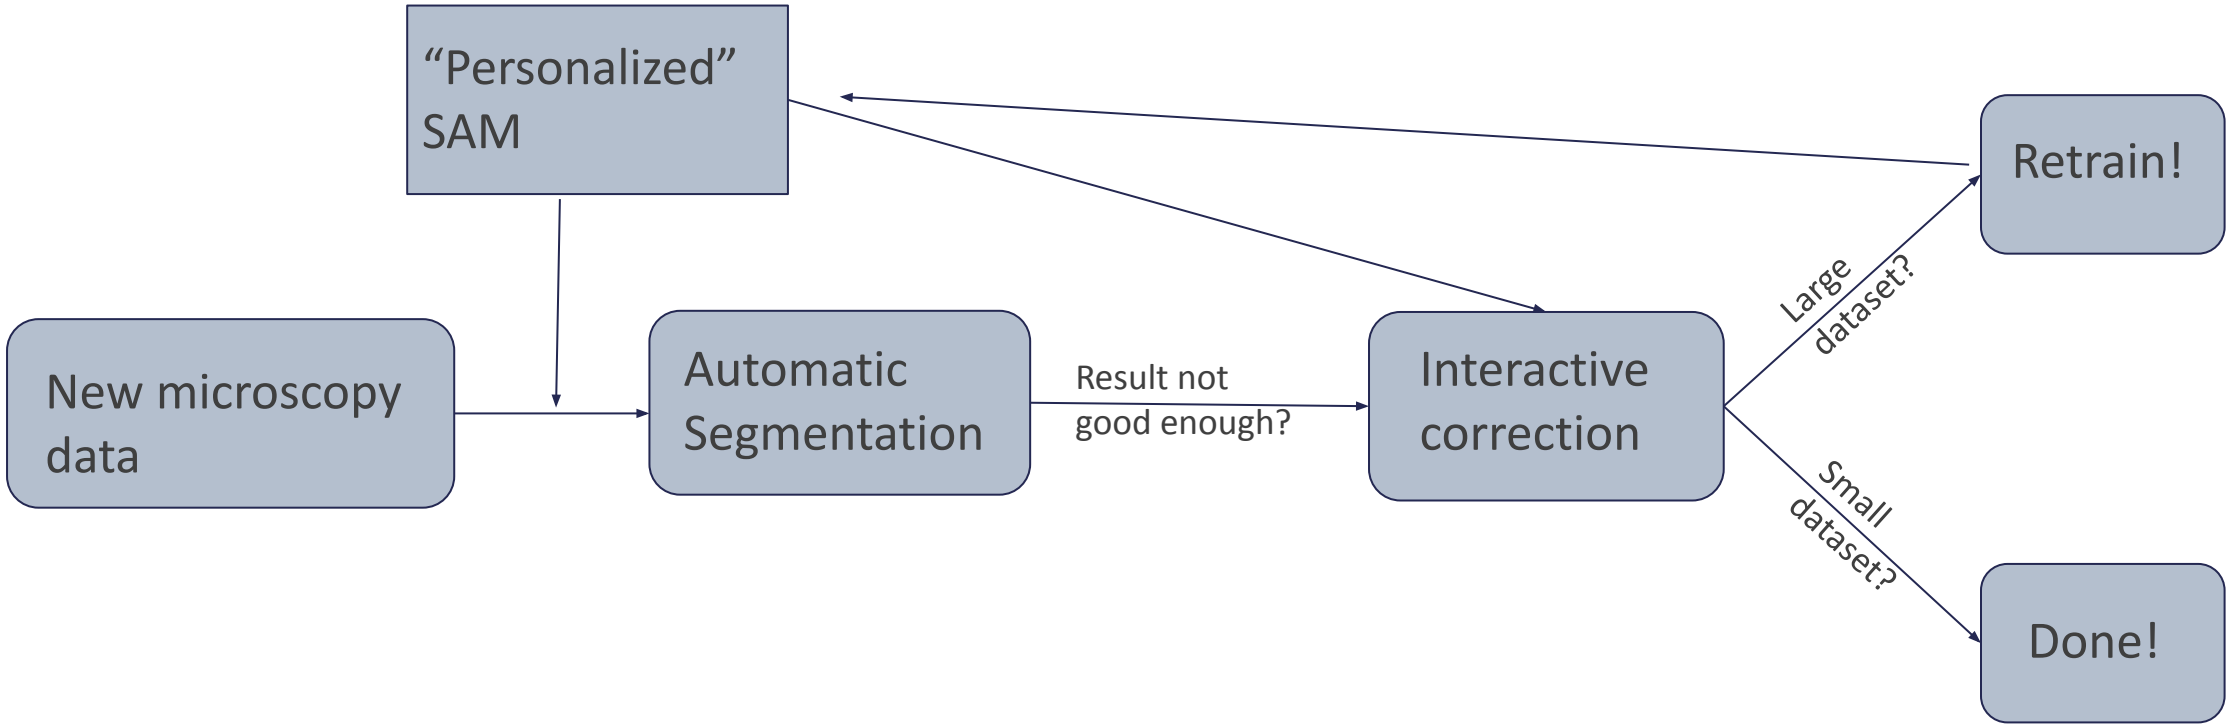

Application in practice

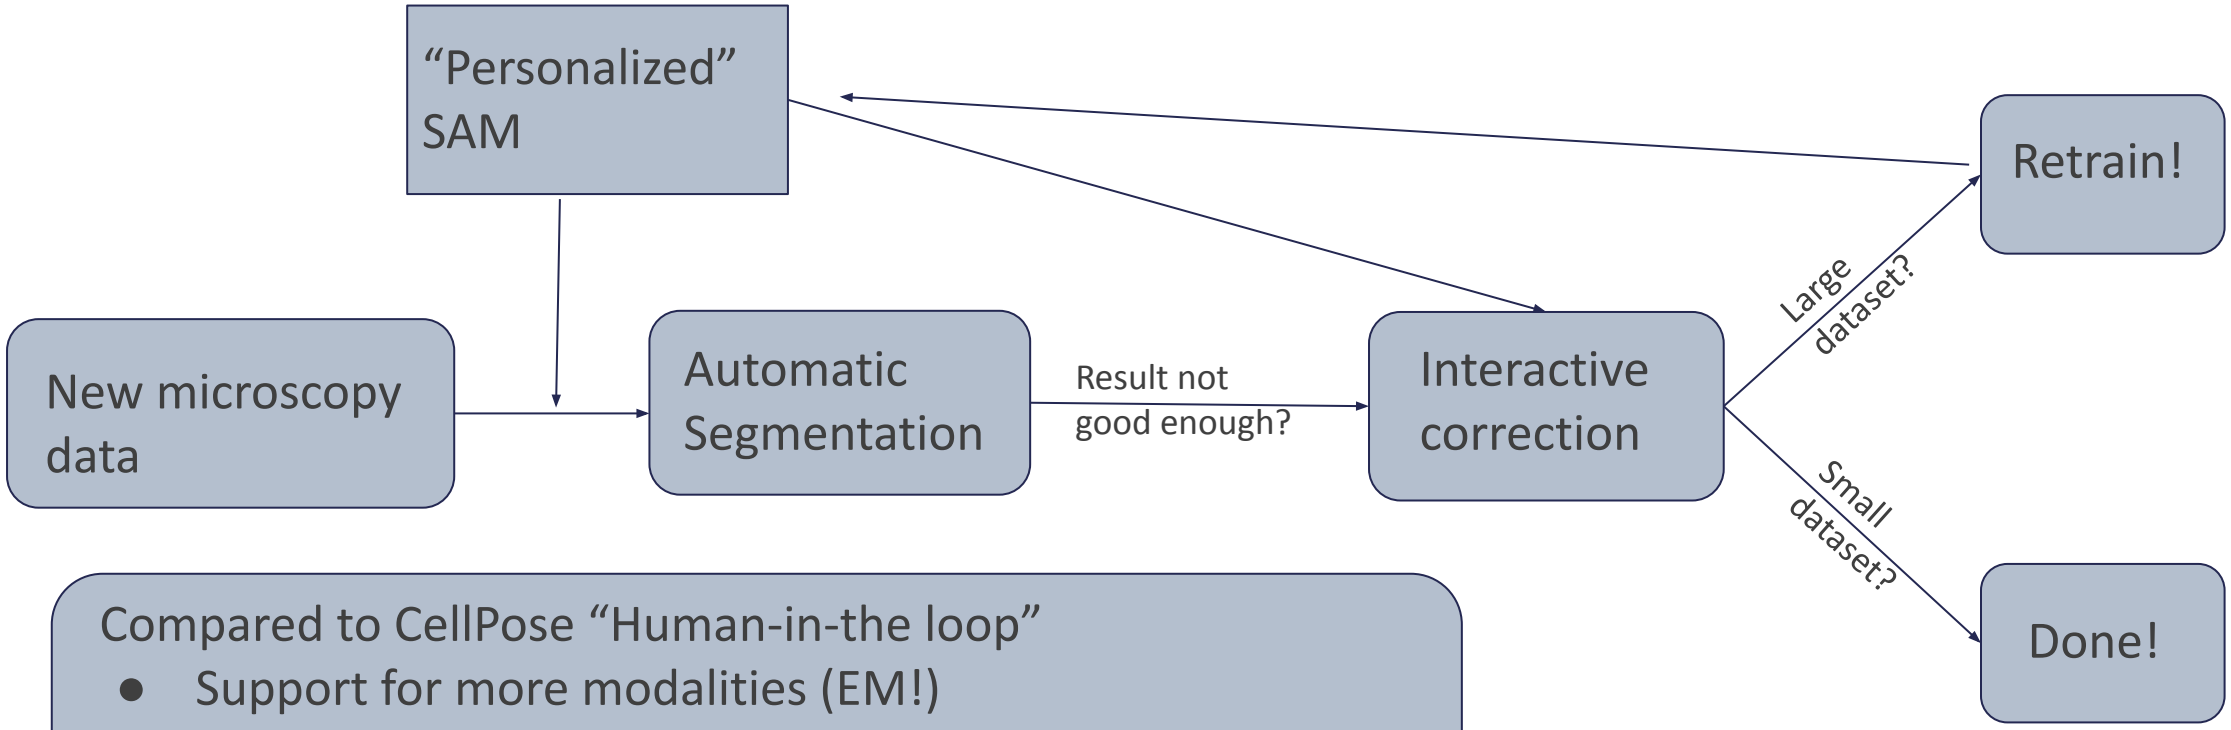

Compared to CellPose "Human-in-the loop"

- Support for more modalities (EM!)
- Interactive correction speeds up annotation significantly!
- **BUT**: Training model takes longer (esp. on CPU)

microSAM: Napari Integration

microSAM: SAM for napari

- napari plugins that enable interactive and automatic:
  - 2D Segmentation
  - 3D Segmentation
  - Tracking (2D + time)
  - Finetuning on own data

- Core functionality:
  - Default + finetuned models
  - Multidimensional segmentation / tracking (interactive and automatic)
  - Tiled prediction for large images

microSAM: SAM for napari

- napari plugins that enable interactive and automatic:
  - 2D Segmentation
  - 3D Segmentation
  - Tracking (2D + time)
  - Finetuning on own data

- Core functionality:
  - Default + finetuned models
  - Multidimensional segmentation / tracking (interactive and automatic)
  - Tiled prediction for large images

Code and documentation available at:
https://github.com/computational-cell-analytics/micro-sam

New release (v0.5):

- Latest microscopy models, compatible with BioImage.IO modelzoo.
- Updated and extended UI, napari plugin integration.
- Will be announced later this or early next week (it's done, but we need to test it and update documentation).

Plan Live Demos

- Starting the tool, explain components
- 2D Segmentation on LiveCELL
  - Compare default and finetuned model (vit_b, show auto segmentation for vit_b)
- 2D Segmentation with tiling (with vit_t)
- 3D Segmentation on Lucchi
  - Use precomputed embeds and amg
- Finetuning (on the Lucchi data we have annotated)

Next Steps & Outlook

Feedback and contributions on the tool are very welcome!

Next steps

- Create v1.0 release: same as v0.5 with additional:
  - Full BioImage.IO integration to enable cross-compatibility.
    - Microscopy Image Browser, QuPath, BioEngine, …
- Integration of efficient training procedures for finetuning (LoRA)
  - To enable better training on CPU and small GPUs
- Provide better and more models:
  - EM Organelle Generalist Model
    - Training on OpenOrganelle and other organelle segmentation datasets.
  - Histopathology Model

Check out our repository for all the details:
https://github.com/computational-cell-analytics/micro-sam

Outlook:
Universal microscopy segmentation and tracking

- Incorporate 3D (2D + time) segmentation in SAM-like model
  - Advantage Transformer: same model for 2d and 3d is possible!
- Vision Mamba: Investigate newer (more efficient) architectures
  - Our recent (preliminary!) work: https://arxiv.org/abs/2404.07705
- Semantic awareness (e.g. differentiate organelles in EM, one model for microscopy)
- Zero-shot adaptation (improve segmentation from examples)

Acknowledgments

EMBL Heidelberg
Anna Kreshuk & her group
et al.

Uni Göttingen & Campus
Alexander Ecker
Tobias Moser
Silvio Rizzoli
et al.

My group
Anwai Archit
Luca Freckmann
Sushmita Nair
Marei Freitag
Sagnik Gupta
et al.

Göttingen Campus

DFG

SFB 1286

Multiscale Cluster of Bioimaging Excellence
From Molecular Machines to Networks of Excitable Cells

CIDAS Campus-Institut Data Science

SARTORIUS