Medical CT at Extreme Scale using Synchrotrons

Professor Anders Bjorholm Dahl
Technical University of Denmark

August, 2024

Tafforeau / ESRF 2022 UCL, CZI, LADAF

Human Organ Atlas Project

Tafforeau, Walsh, Lee, et al. ESRF

Human Organ Atlas Project

Human Organ Atlas EXPLORE SEARCH 3D RECONSTRUCTIONS TUTORIALS HELP

Welcome to the Human Organ Atlas

The Human Organ Atlas uses **Hierarchical Phase-Contrast Tomography** to span a previously poorly explored scale in our understanding of human anatomy, the micron to whole intact organ scale.

Histology using optical and electron microscopy images cells and other structures with sub-micron accuracy but only on small biopsies of tissue from an organ, while clinical CT and MRI scans can image whole organs, but with a resolution only down to just below a millimetre. HiP-CT bridges these scales in 3D, imaging intact organs with ca. 20 micron voxels, and locally down to microns.

We hope this open access Atlas, enabled by the ESRF-EBS, will act as a reference to provide new insights into our biological makeup in health and disease. To stay up to date, follow @HiP-CT

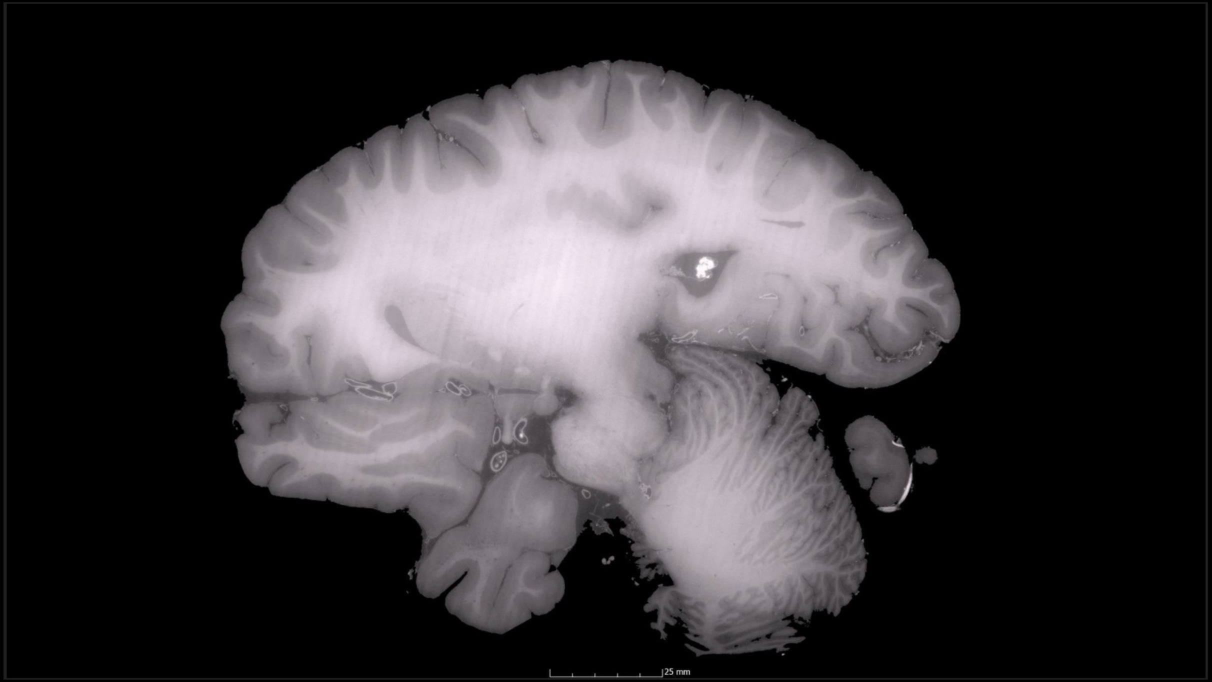

HiP-CT imaging and 3D reconstruction of a complete brain from the body donor LADAF-2020-31. More videos can be viewed on the HiP-CT YouTube channel.

https://human-organ-atlas.esrf.eu/

Organs

Heart

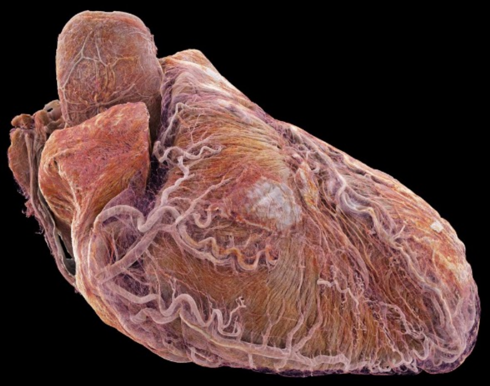

Brain

Left kidney

Right kidney

Spleen

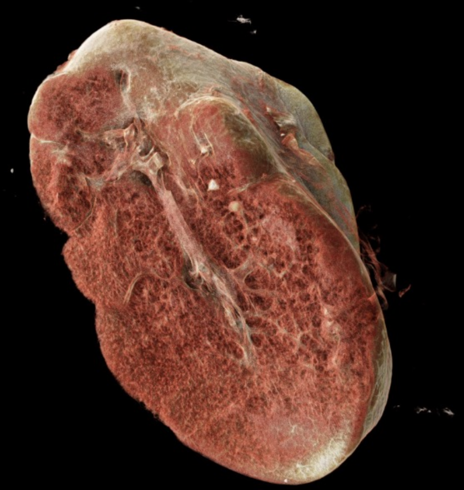

Multiple organs
Multiple scales
Large volumes
Available online

AI for the Human Organ Atlas

Synchrotron CT – from whole organ to nano-meter resolution

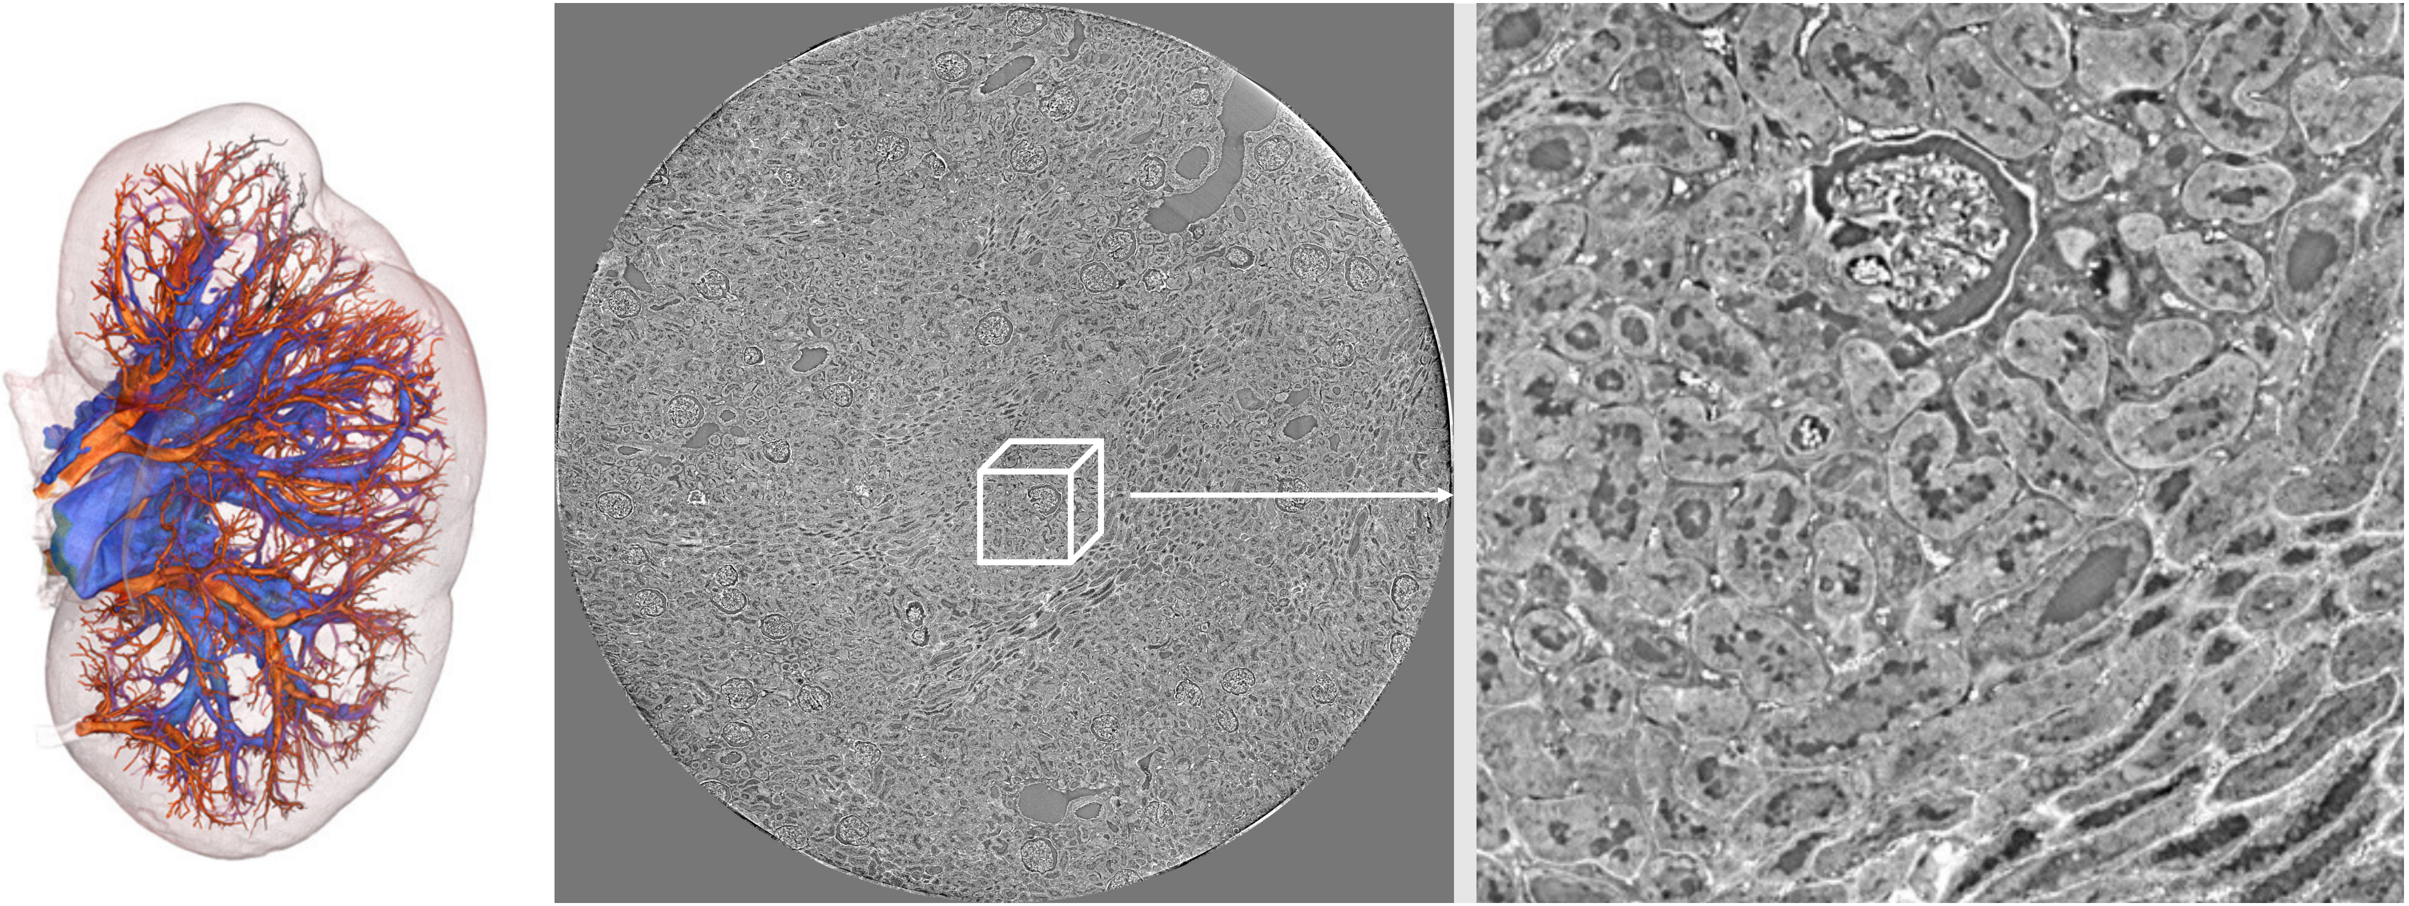

Schmidt-Christensen et al. 2023

Example – optimizing devices

Rat kidney

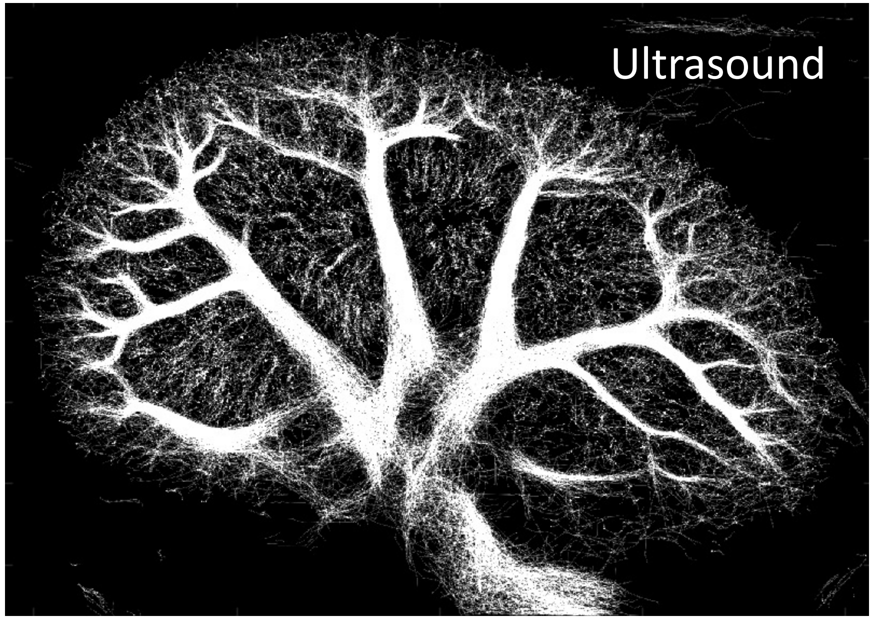

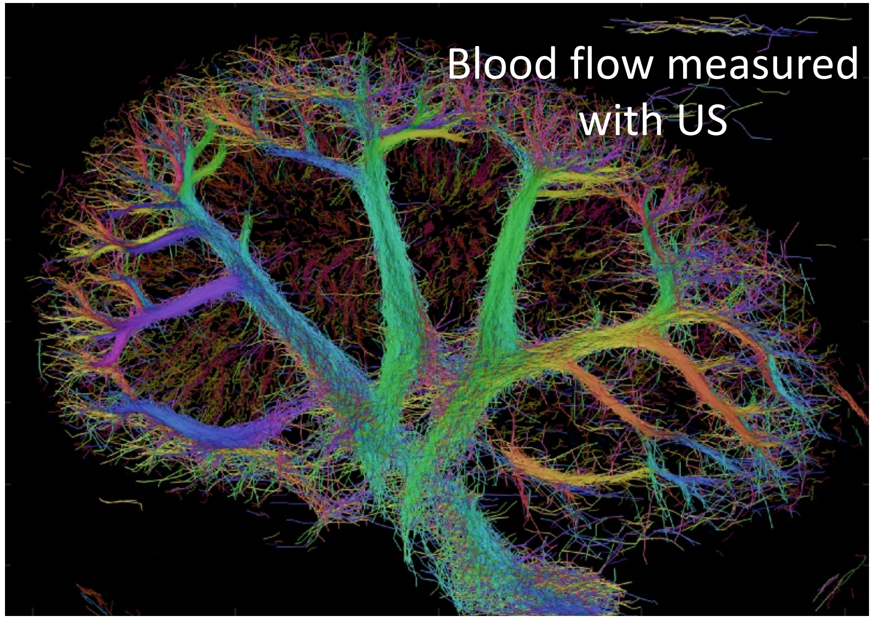

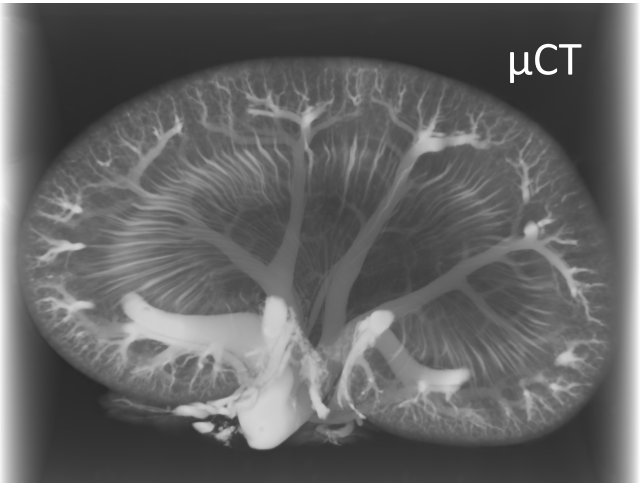

Interactive segmentation of μCT

Imaging pipeline

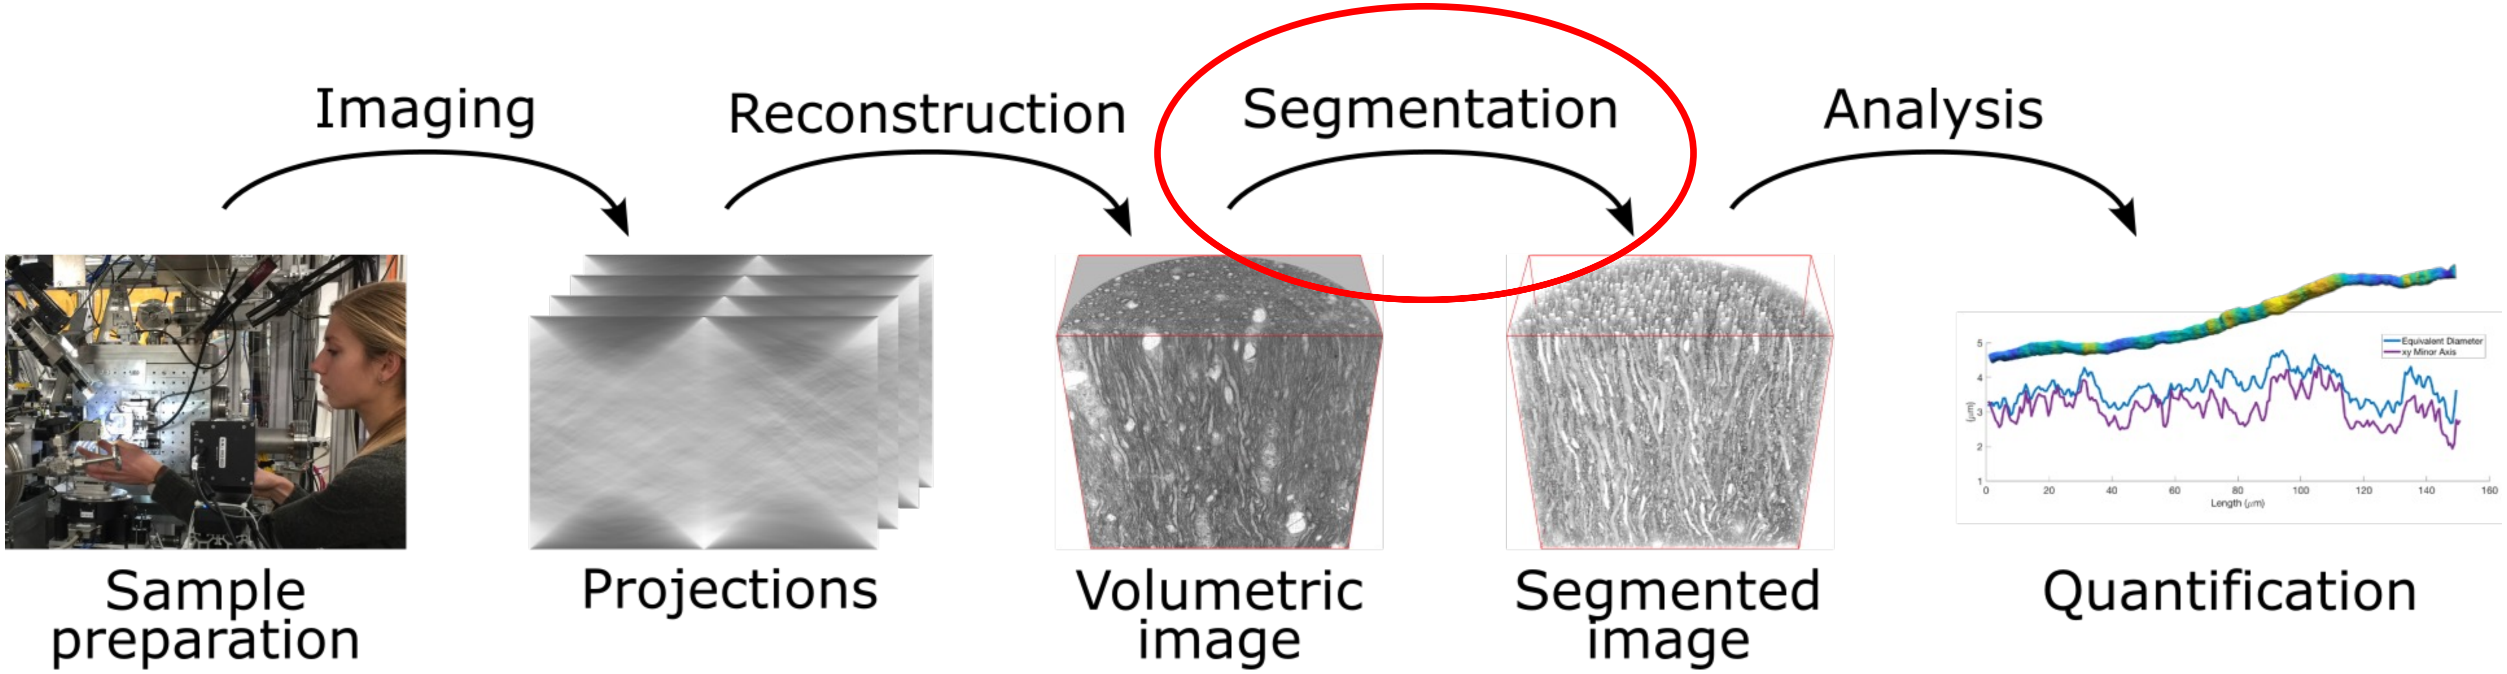

Image segmentation – manual decision

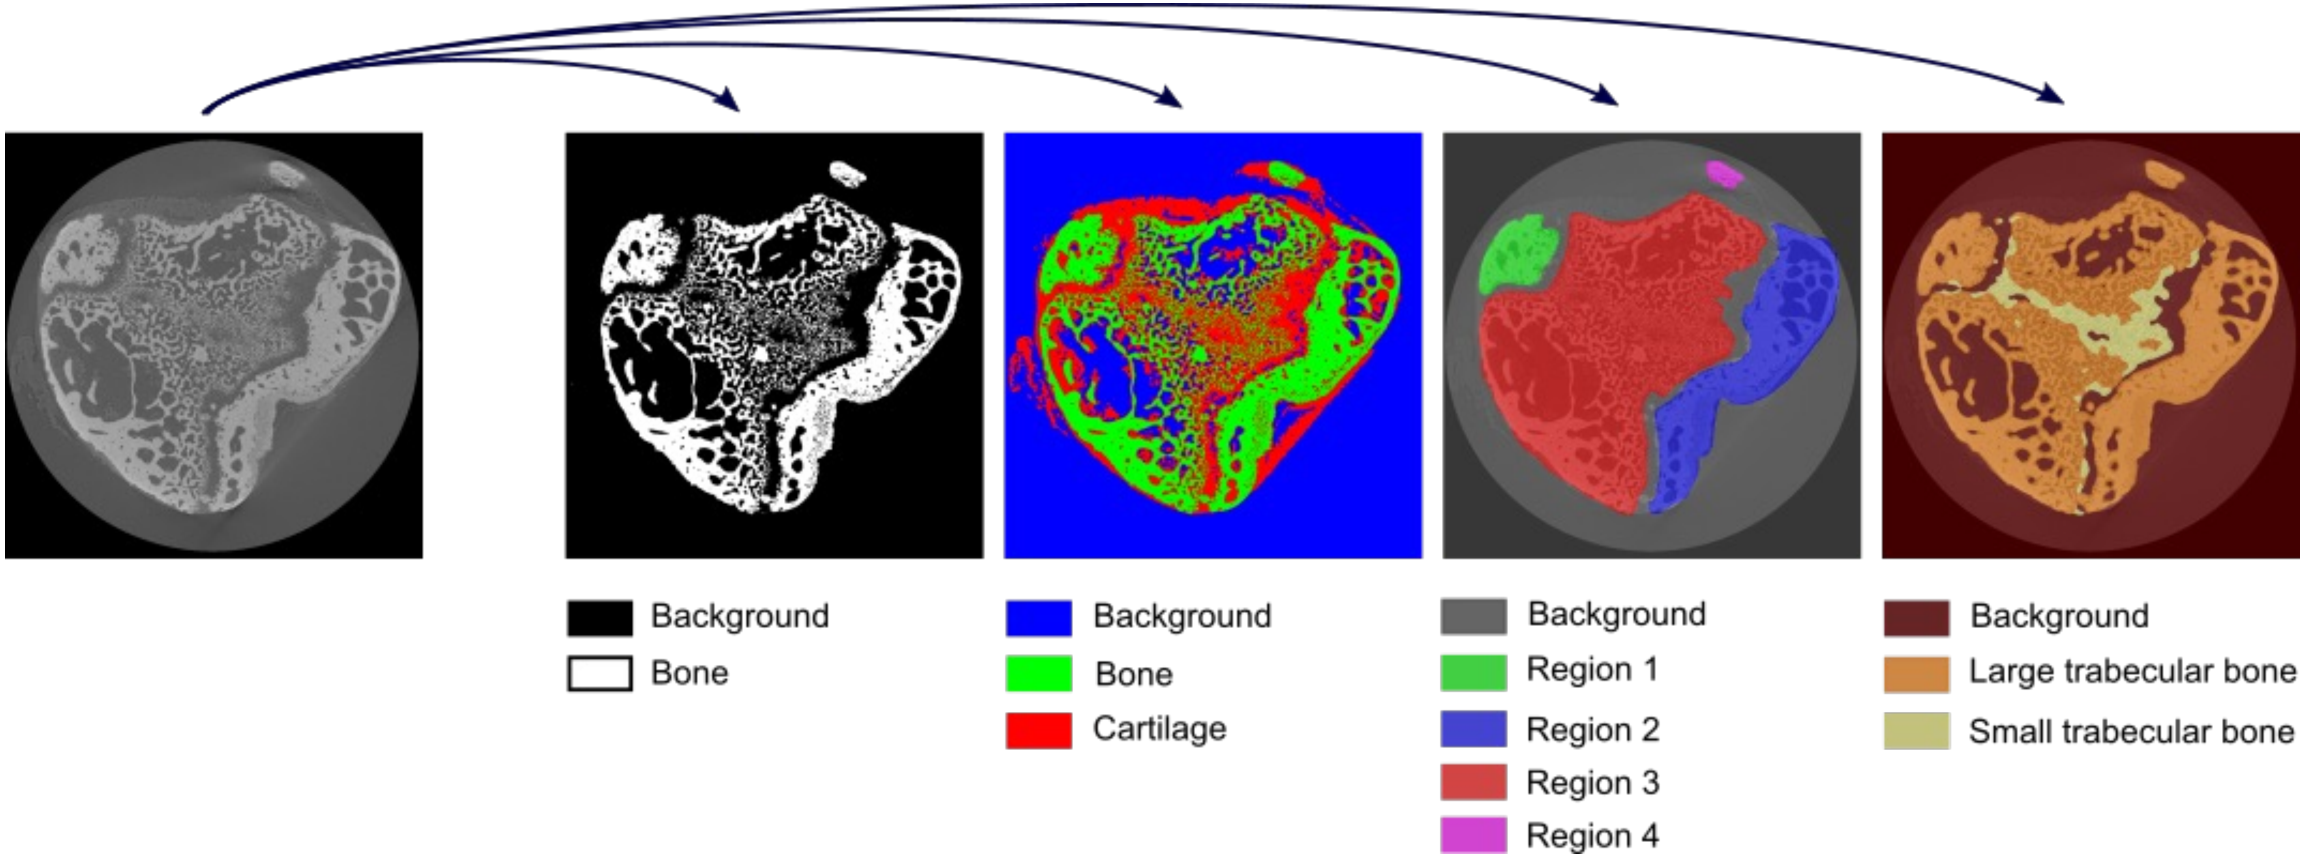

Towards Large Foundational Models

SAM – Segment Anything Model

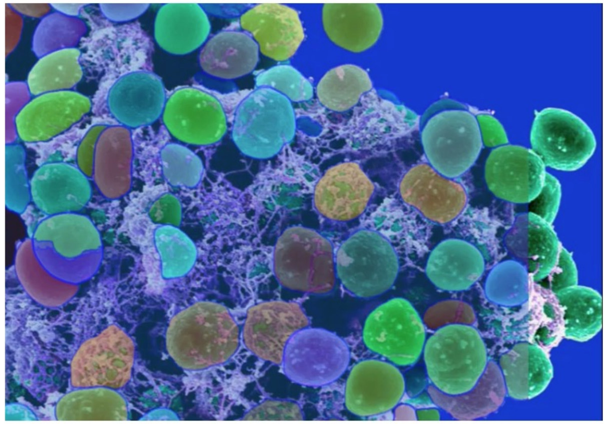  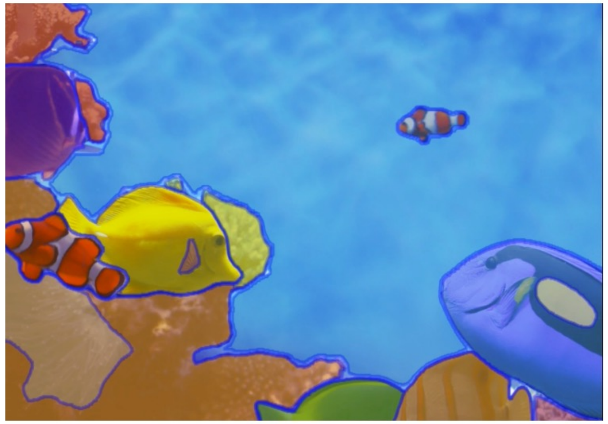

Foundational model

Trained on 11 million images and more than 1 billion labels

Source: https://segment-anything.com/demo
Kirillov, Alexander, et al., ICCV, 2023

Foundational Medical Models

MedSAM – Segment Anything in Medical Images

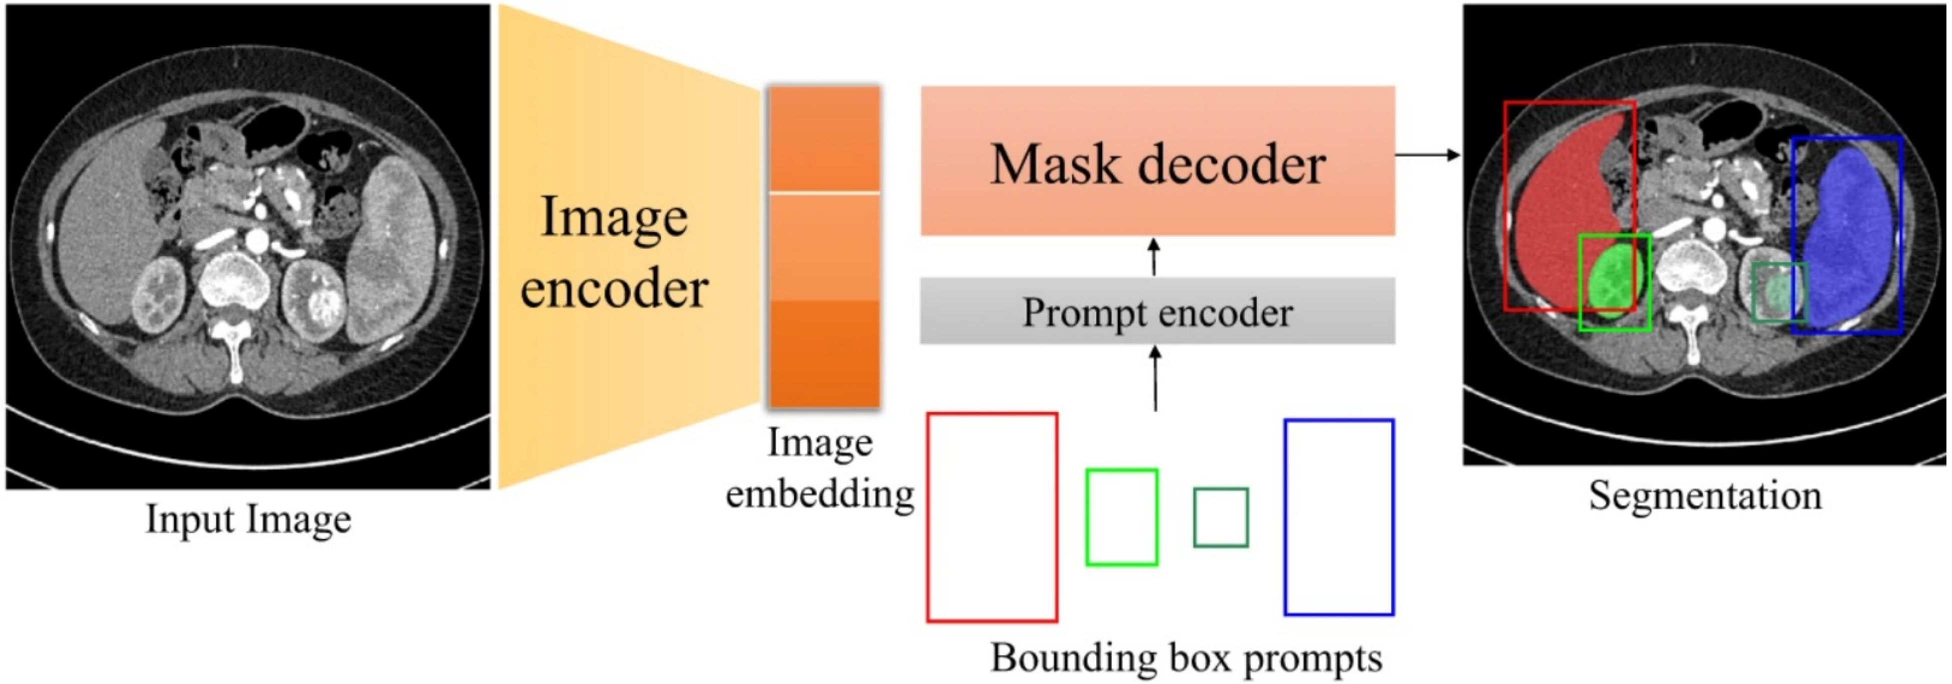

1,570,263 images – 10 imaging modalities
CT, MRI, endoscopy, histology, etc.

Source: Ma, Jun, et al. Nature Communications, 2024

TotalSegmentator

If used for research purposes, please cite our Radiology AI paper.
The results of the models appendicular bones, tissue types, heartchambers highres and face may not be used commercially.
All other results are open for any usage.
If this website does not work, please create a issue on github.

Please help us continue providing TotalSegmentator.com by sharing your thoughts through this quick form. Thank you!

Download results

Upload a new case

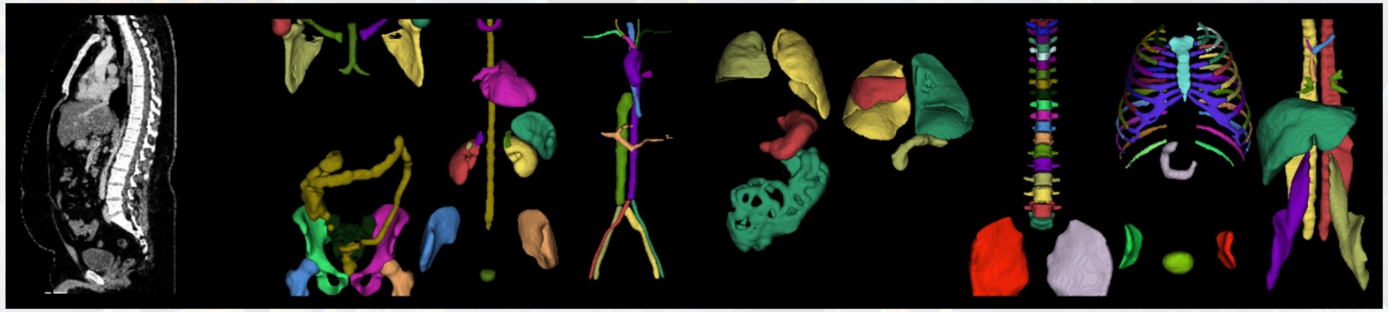

Preview rendering of all segmentations

#	Segmentation Mask (nifti)
1	adrenal_gland_left.nii.gz
2	adrenal_gland_right.nii.gz
3	aorta.nii.gz
4	atrial_appendage_left.nii.gz
5	autochthon_left.nii.gz

Facts: Trained from 1240 annotated CT scans – 104 anatomical structures

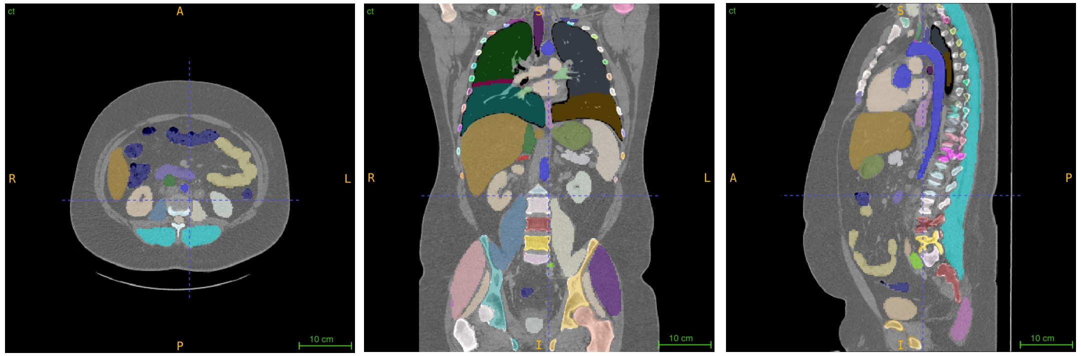

Ground truth

Segmentation

TotalSegmentator v2.0.0

Try out the TotalSegmentator by uploading any CT data. The upload must meet the following criteria:

- Only a single CT dataset, maximum size is 400 MB
- Upload should be either a zip file of DICOMs or a single NIFTI image

By using this online service you agree that the data can be used to improve the model.

Drop **DICOM.zip** or **NIFTI.nii.gz** here
or click to upload

Selected task
total (default)

Choose a subset to avoid long runtime (if none is selected all are selected)

Enable fast processing

Calculate statistics (volume and intensity)

Process data

If used for research purposes, please cite our Radiology AI paper.
The results of the models appendicular bones, tissue types, heartchambers highres and face may not be used commercially.
All other results are open for any usage.
If this website does not work, please create a issue on github.

Segmentation

Radiology: Artificial Intelligence

ORIGINAL RESEARCH

TotalSegmentator: Robust Segmentation of 104 Anatomic Structures in CT Images

Jakob Wasserthal, PhD • Hanns-Christian Breit, MD • Manfred T. Meyer, MD • Maurice Pradella, MD • Daniel Hinck • Alexander W. Sauter, MD • Tobias Heye, MD • Daniel T. Boll, MD • Joshy Cyriac, MSc • Shan Yang, PhD • Michael Bach, PhD • Martin Segeroth, MD

From the Clinic of Radiology and Nuclear Medicine, University Hospital Basel, Basel, Switzerland, Petersgraben 4, 4031 Basel, Switzerland. Received January 25, 2023; revision requested February 27; revision received May 16; accepted June 14. **Address correspondence to** J.W. (email: *jakob.wasserthal@usb.ch*).

Authors declared no funding for this work.

Conflicts of interest are listed at the end of this article.

See also commentary by Sebro and Mongan in this issue.

Radiology: Artificial Intelligence 2023; 5(5):e230024 • https://doi.org/10.1148/ryai.230024 • Content codes: AI CA CT GI MK

showed a high Dice score (0.943) on t
nificantly outperformed another public

Purpose: To present a deep learning segmentation model that can automatically and robustly segment all major anatom
body CT images.

Materials and Methods: In this retrospective study, 1204 CT examinations (from 2012, 2016, and 2020) were used to segment 104 anatomic structures (27 organs, 59 bones, 10 muscles, and eight vessels) relevant for use cases such as organ volumetry, disease characterization, and surgical or radiation therapy planning. The CT images were randomly sampled from routine clinical studies and thus represent a real-world dataset (different ages, abnormalities, scanners, body parts, sequences, and sites). The authors trained an nnU-Net segmentation algorithm on this dataset and calculated Dice similarity coefficients to evaluate the model's performance. The trained algorithm was applied to a second dataset of 4004 whole-body CT examinations to investigate age-dependent volume and attenuation changes.

Results: The proposed model showed a high Dice score (0.943) on the test set, which included a wide range of clinical data with major abnormalities. The model significantly outperformed another publicly available segmentation model on a separate dataset (Dice score, 0.932 vs 0.871; $P < .001$). The aging study demonstrated significant correlations between age and volume and mean attenuation for a variety of organ groups (eg, age and aortic volume [$r_s = 0.64$; $P < .001$]; age and mean attenuation of the autochthonous dorsal musculature [$r_s = -0.74$; $P < .001$]).

Performance measures

Accuracy:

$$\text{Acc} = \frac{\text{True pixels}}{\text{Total pixels}}$$

Intersection over Union:

$$\text{IoU} = \frac{|X \cap Y|}{|X \cup Y|}$$

Dice:

$$\text{DSC} = \frac{2|X \cap Y|}{|X| \cup |Y|}$$

Performance measures

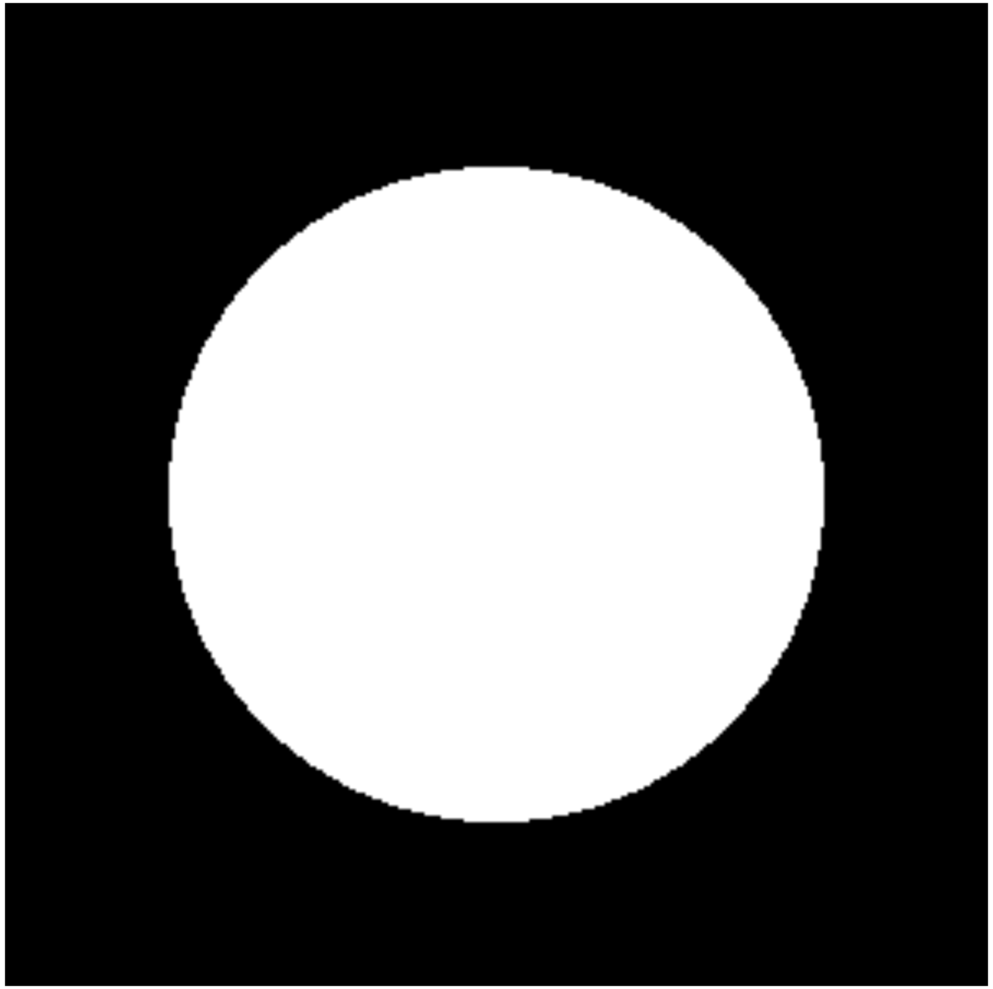

Performance measures

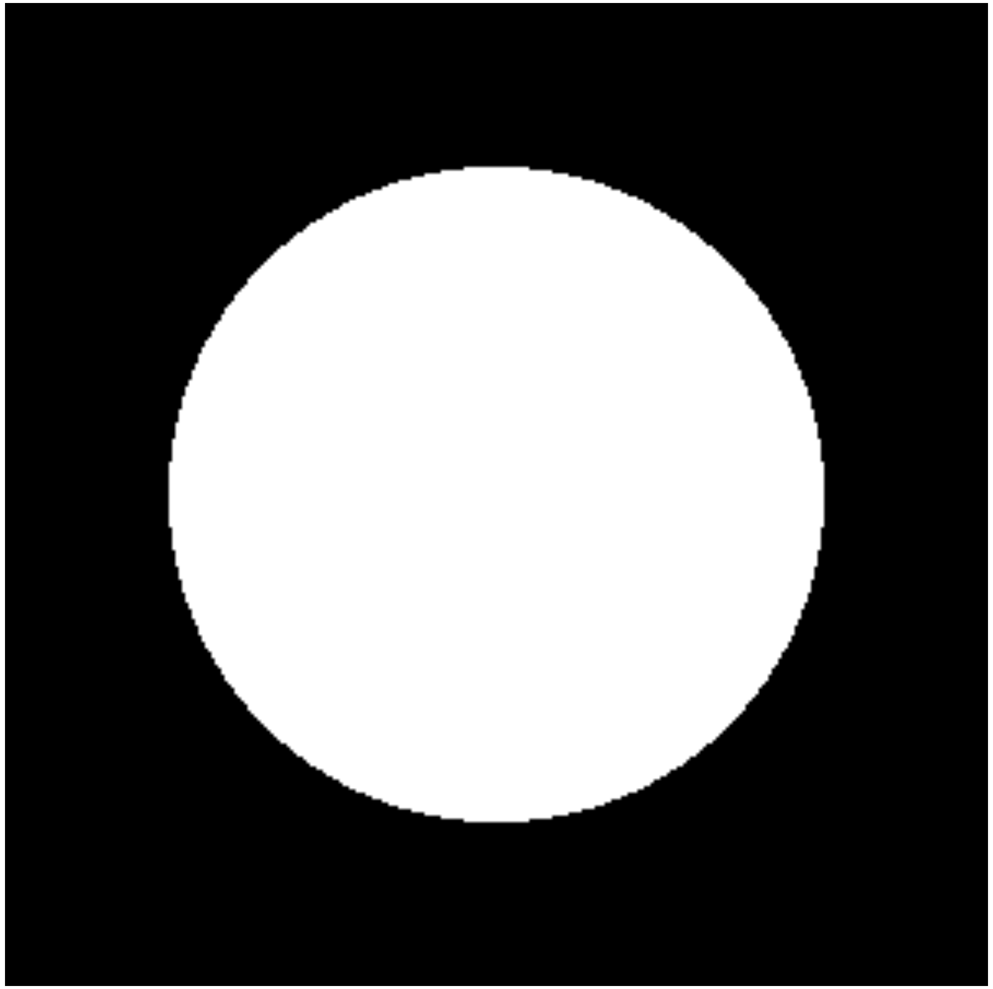

Original

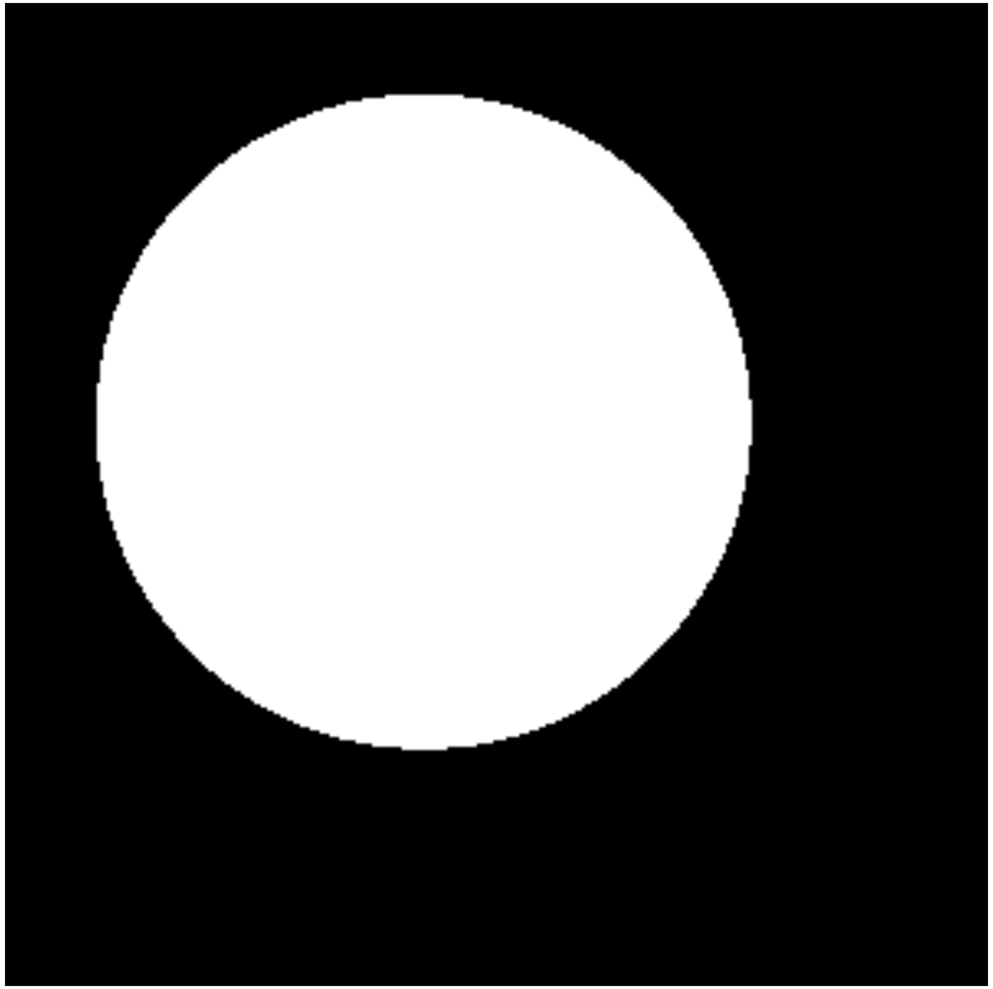

Segmentation

Performance measures

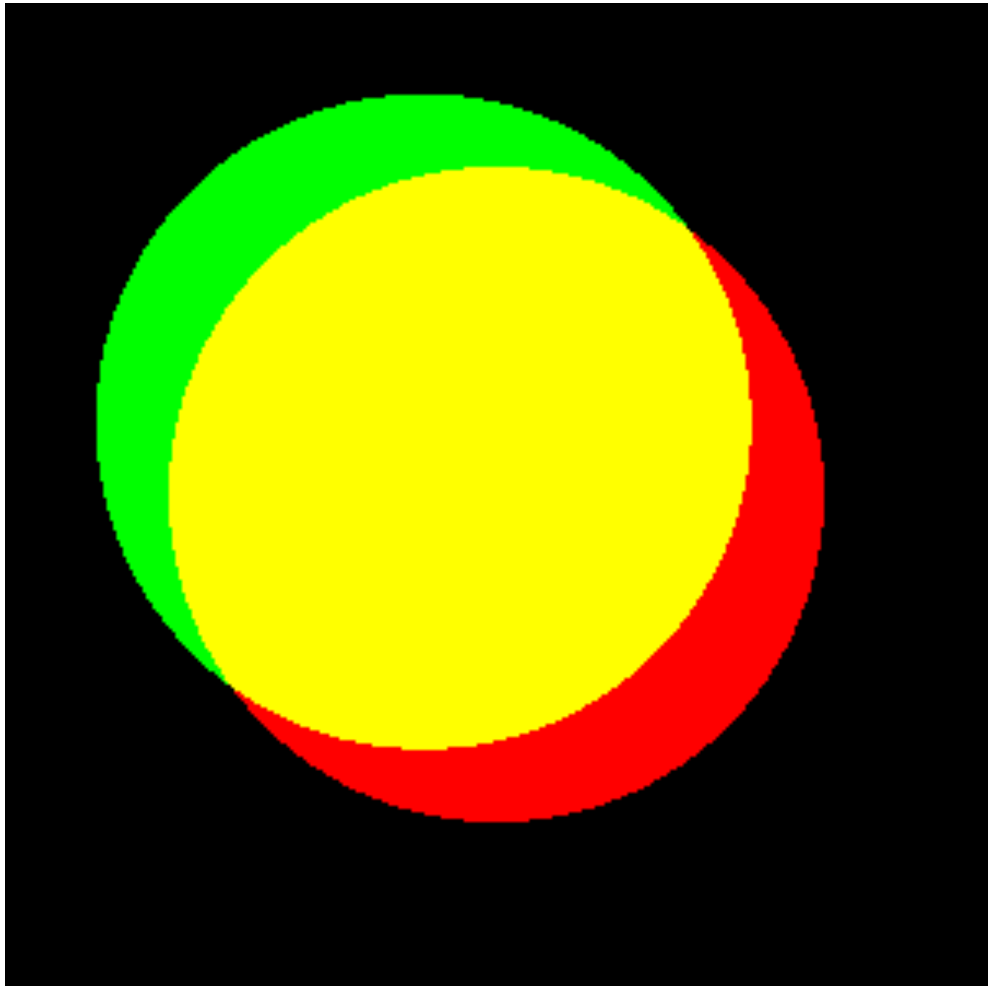

Overlap

Accuracy:	0.861
Intersection over Union:	0.667
Dice:	0.800

Performance measures

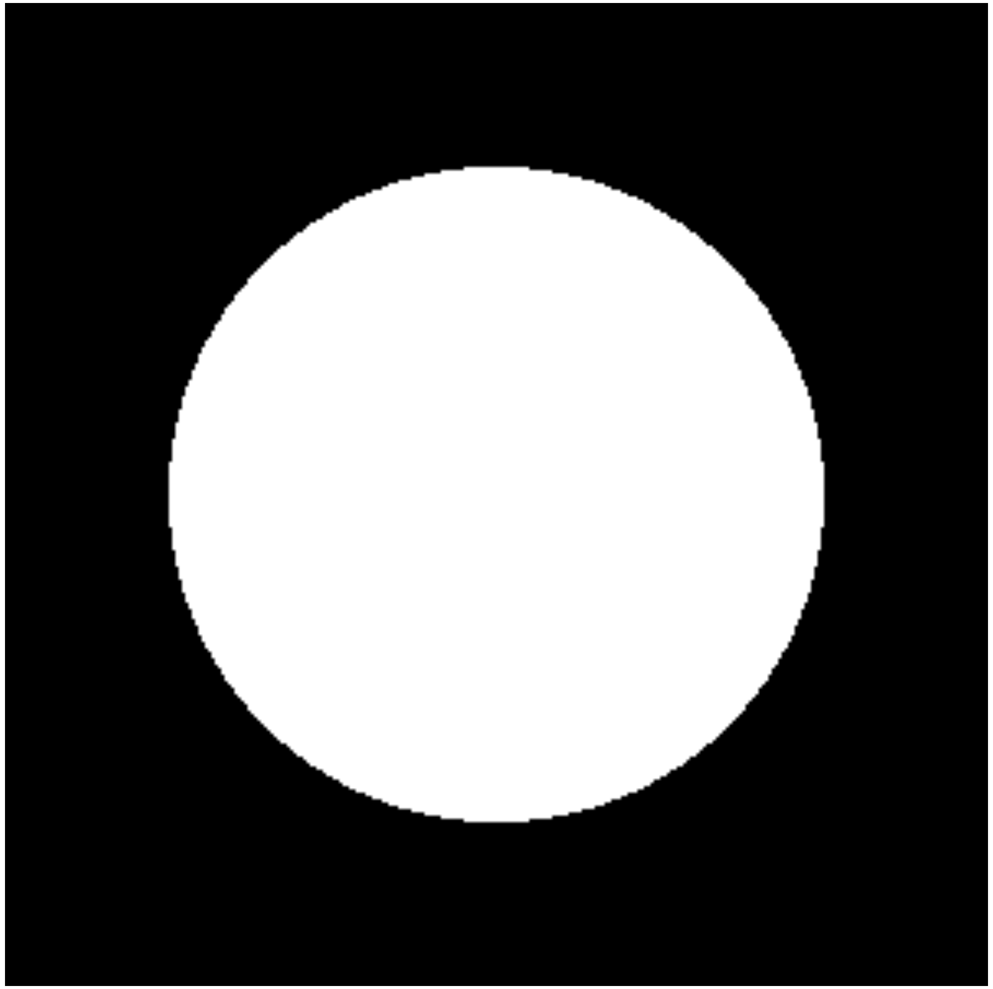

Original

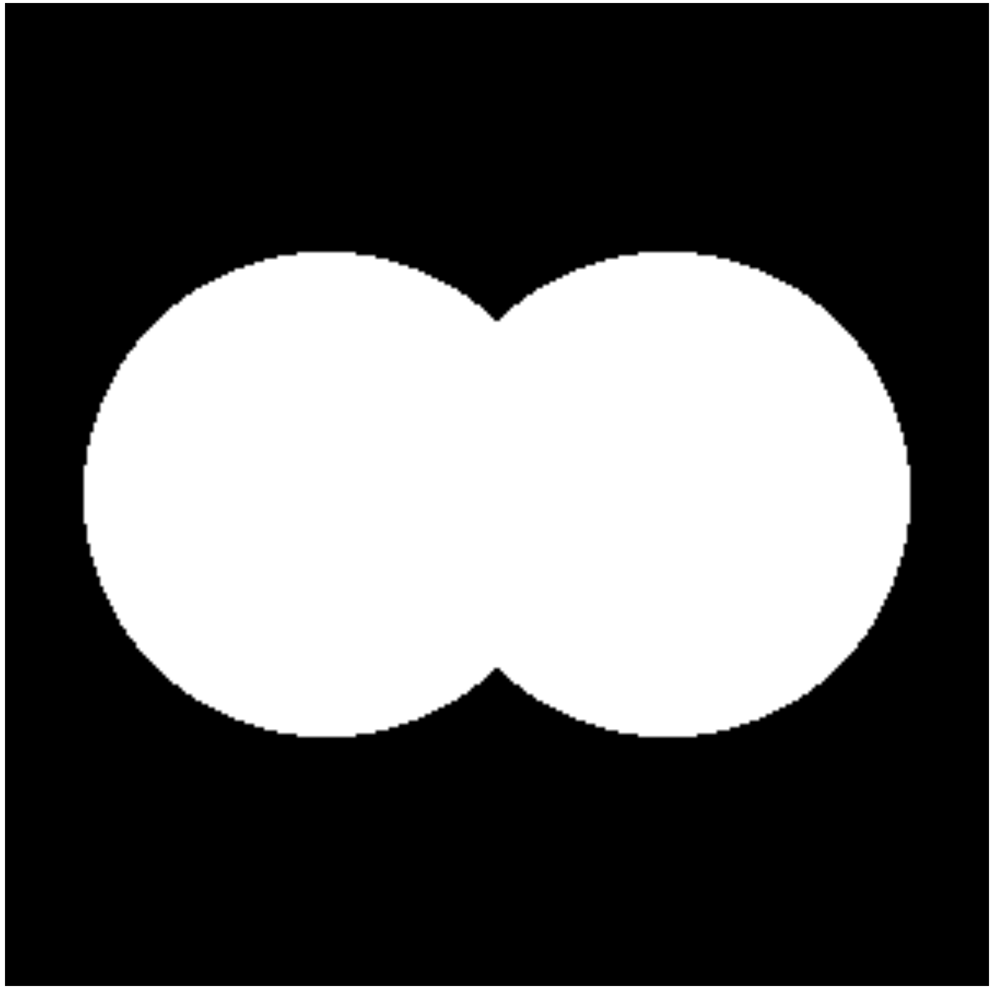

Segmentation

Performance measures

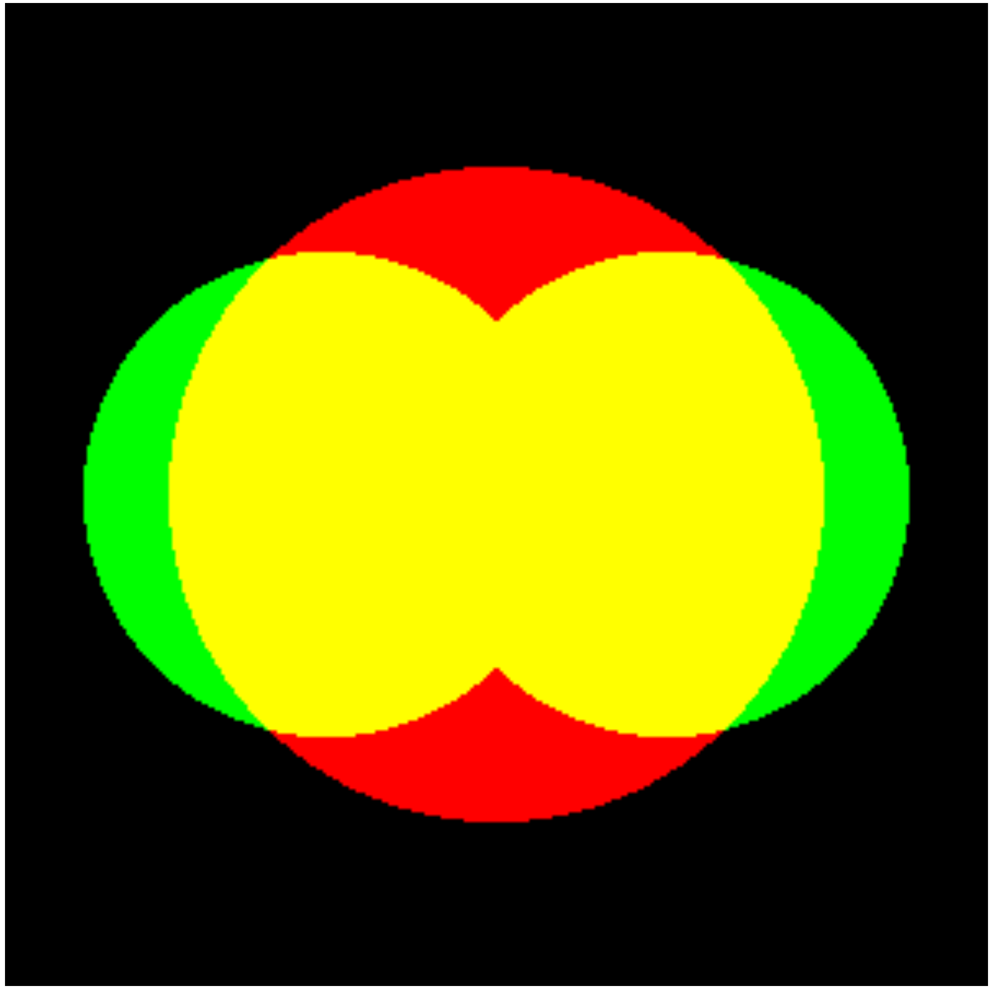

Overlap

Accuracy:	0.864
Intersection over Union:	0.673
Dice:	0.804

Performance measures

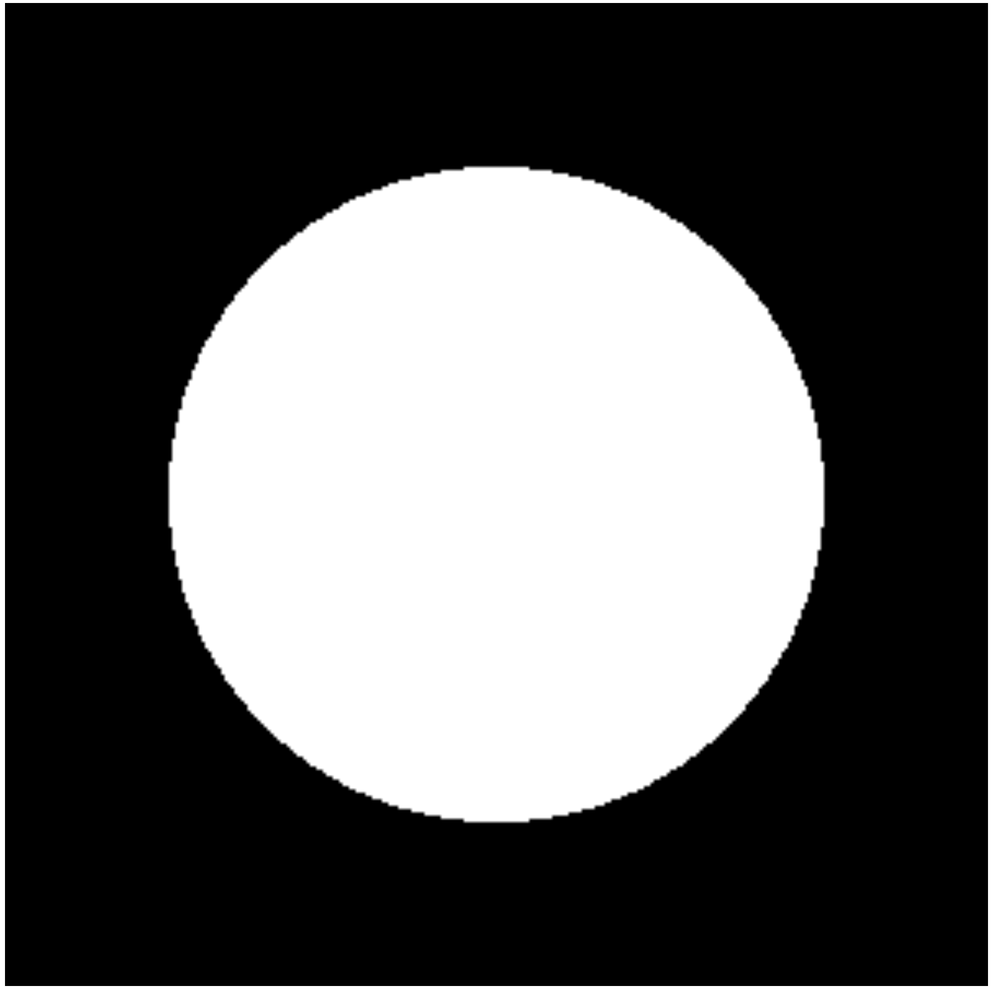

Original

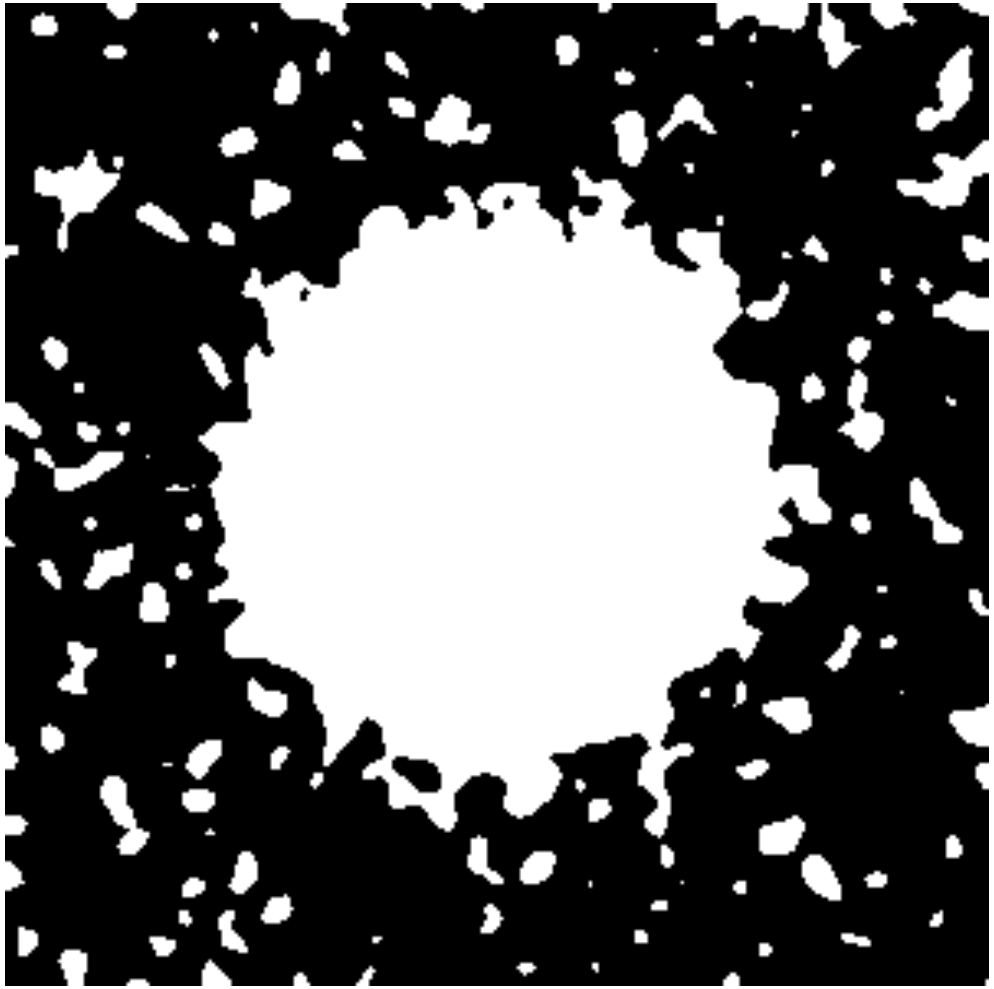

Segmentation

Performance measures

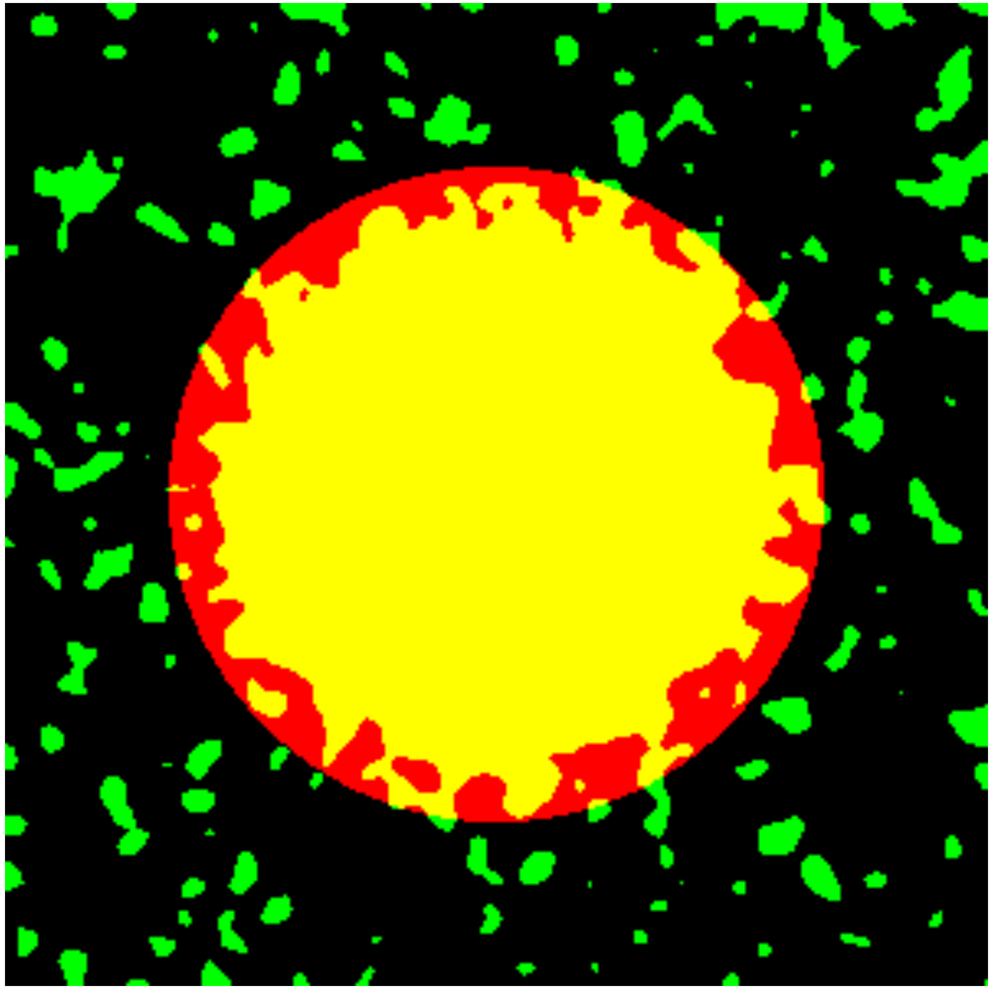

Overlap

Accuracy:	0.861
Intersection over Union:	0.668
Dice:	0.801

Quantitative measures

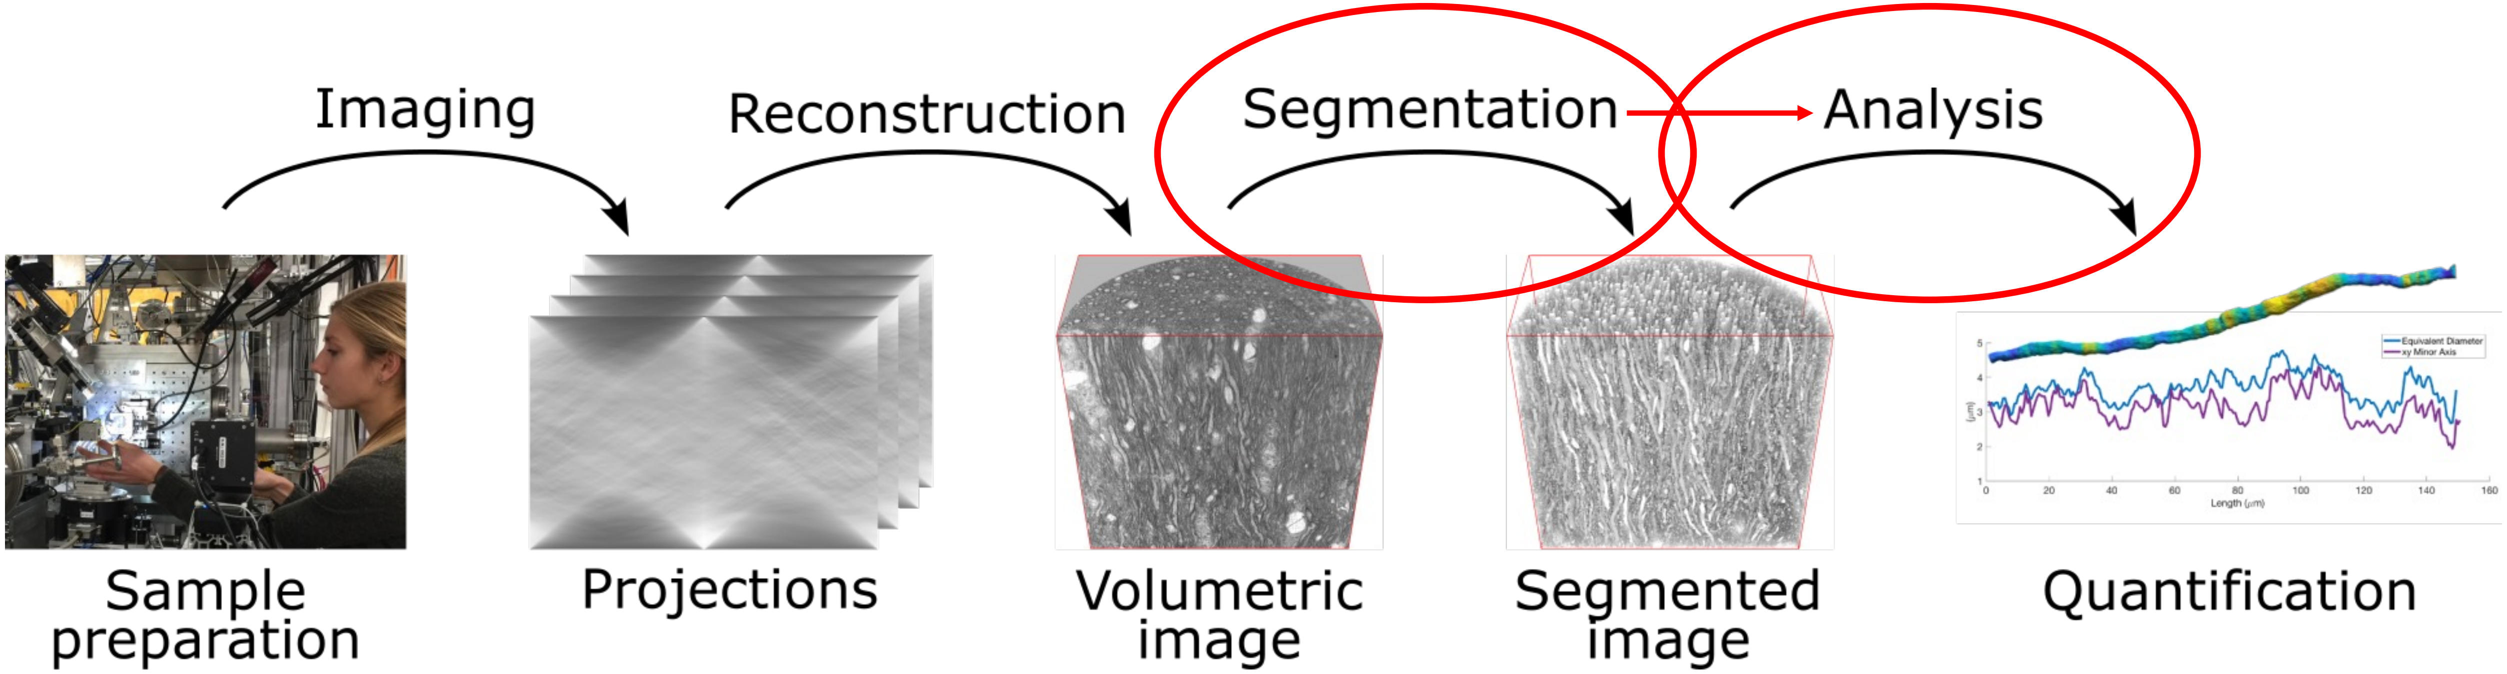

Quantitative measures

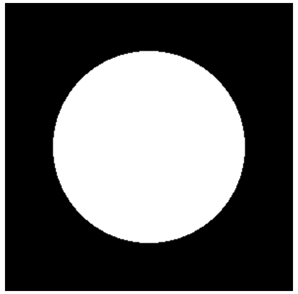 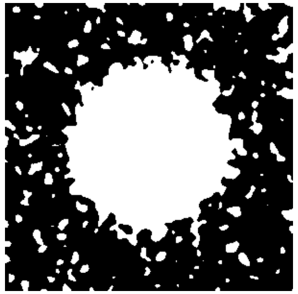

Boundary length:	Short	Long
Number of components:	Few (one)	Many
Local thickness/spacing (median):	High	Low

Local thickness

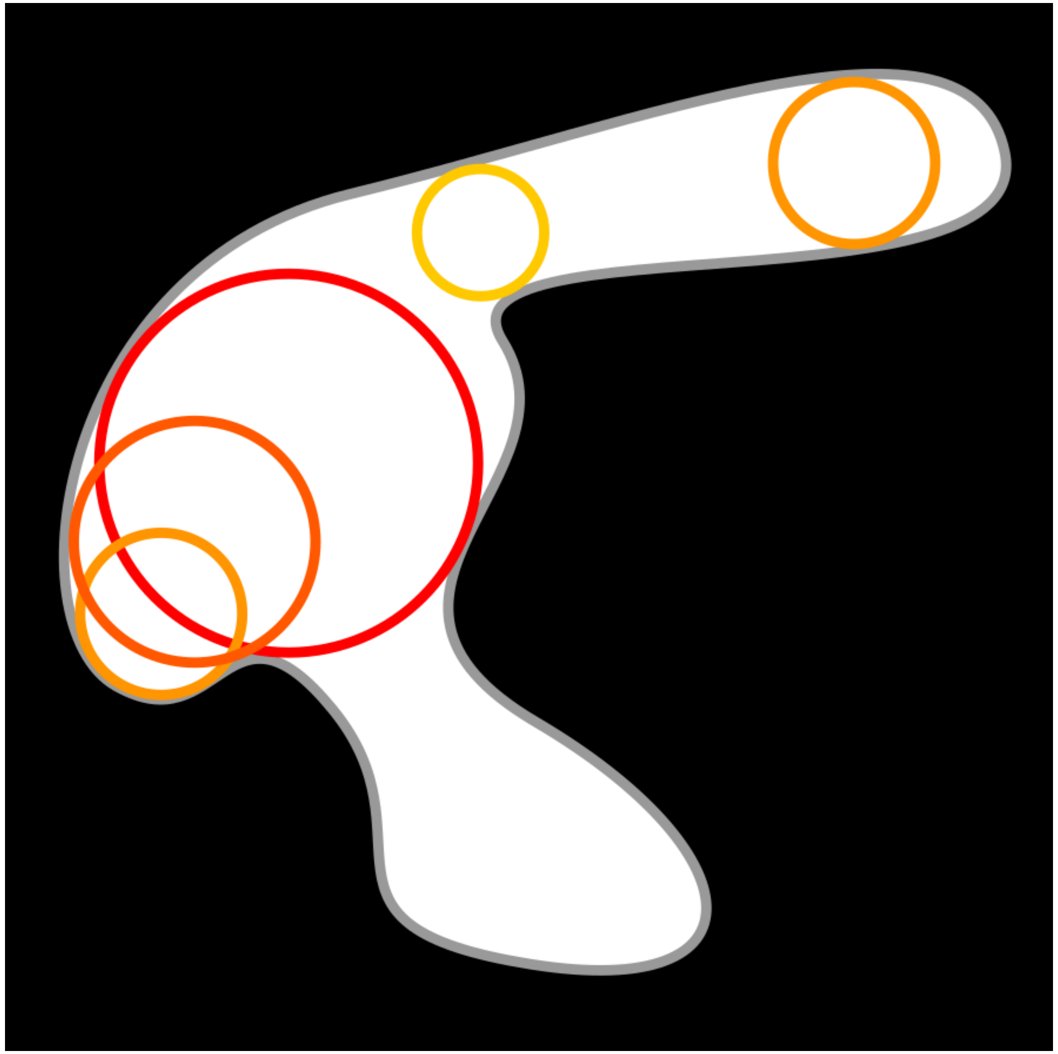

Local Thickness in 3D

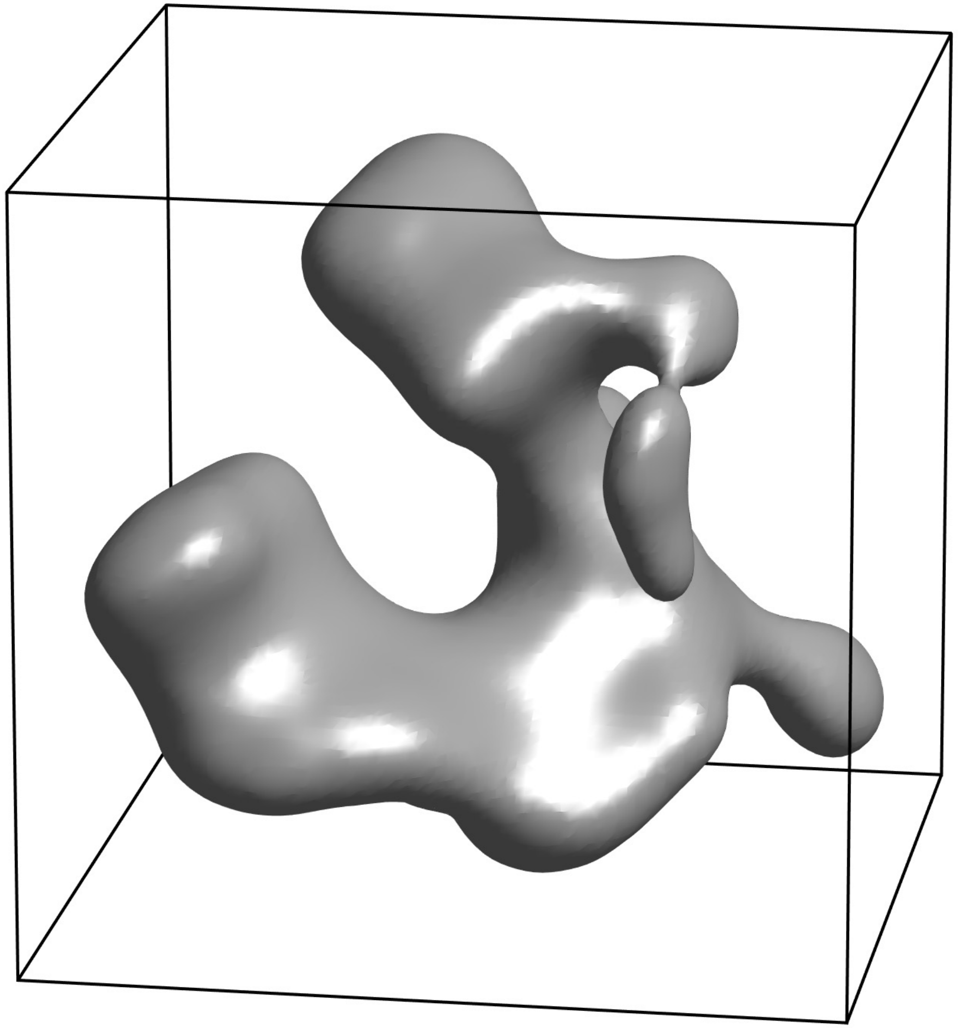

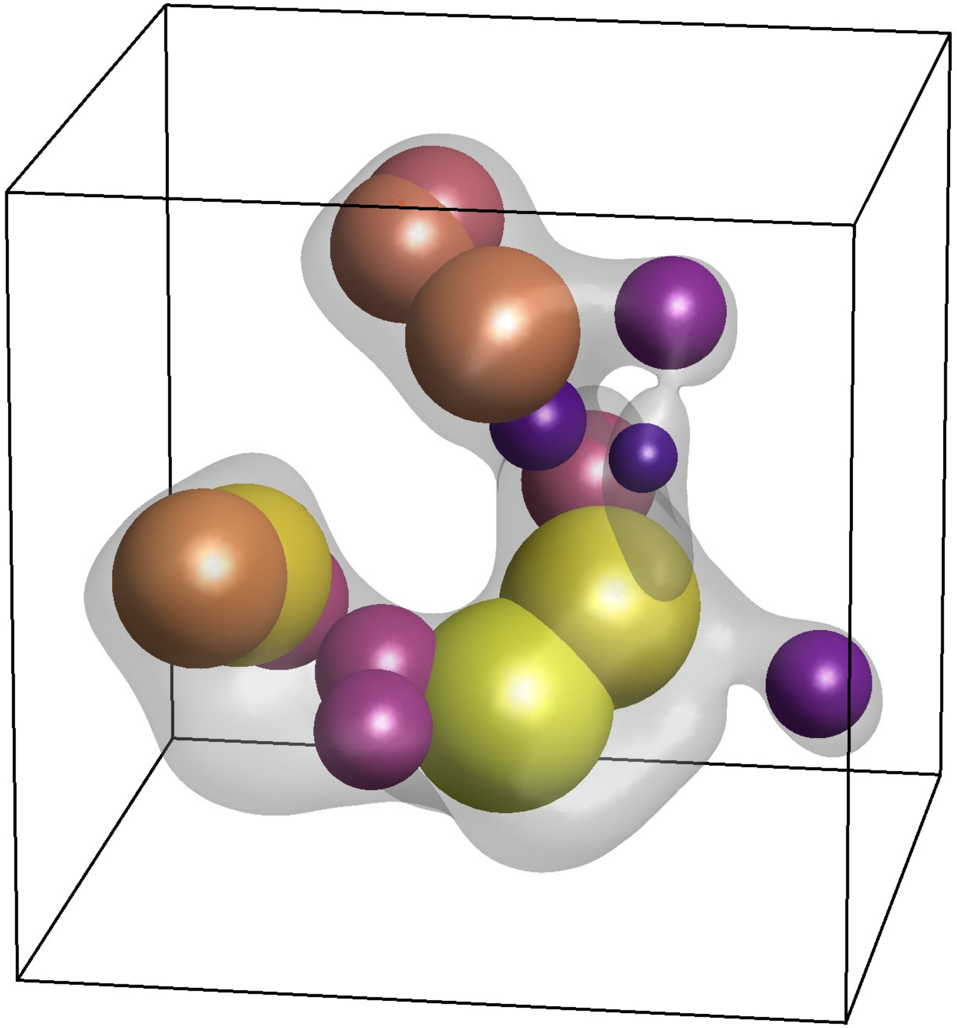

Fast Local Thickness

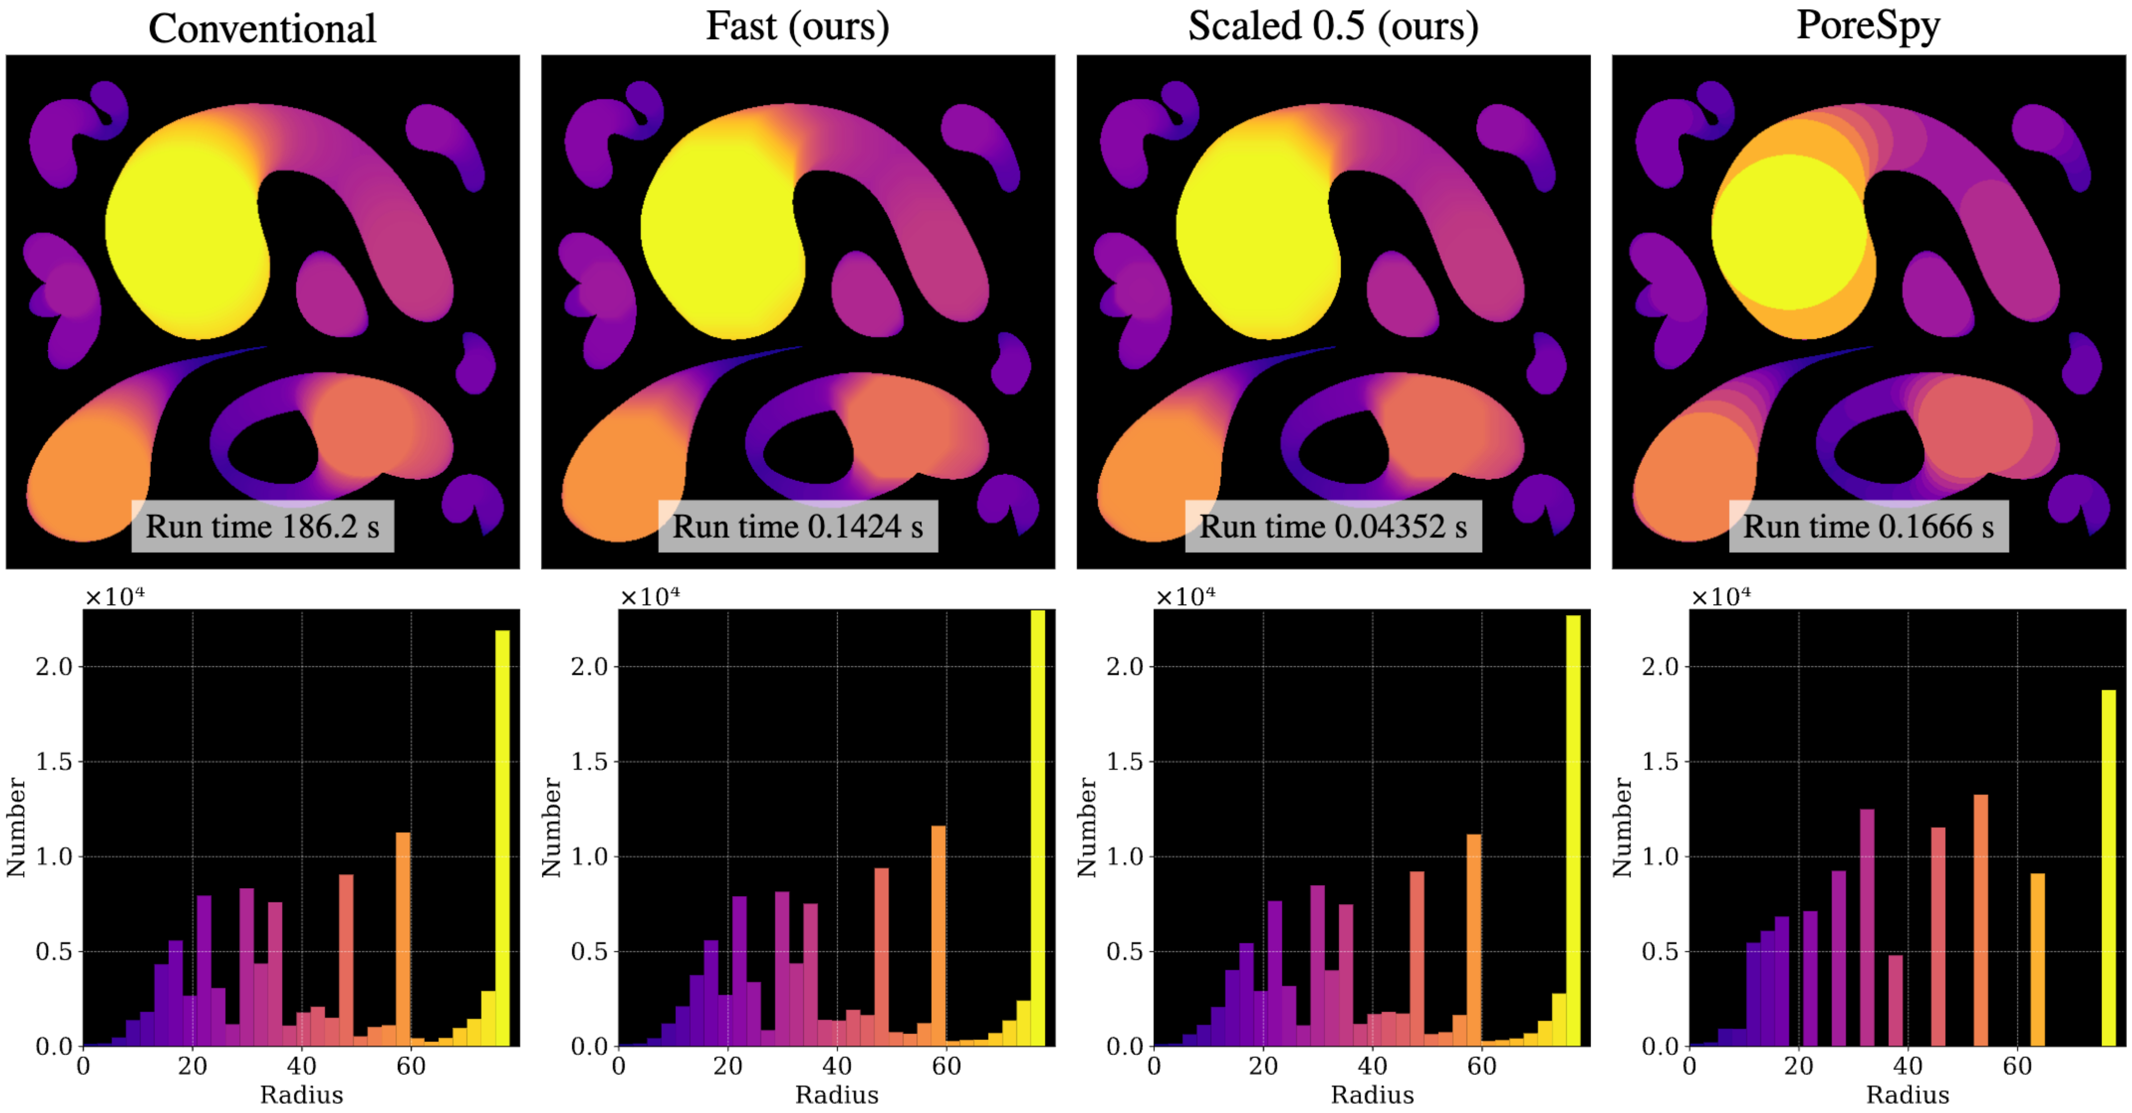

Source: V. Dahl & A. Dahl, 2023, CVPR Workshops

Thickness statistics

Distribution of thickness values:

$$\text{mean} = \frac{1}{n}\sum_{i=1}^{n} h_i$$

$$\mu_{\log} = \frac{1}{n}\sum_{i=1}^{n} \log(h_i)$$

$$V_{\log} = \frac{1}{n}\sum_{i=1}^{n} (\log(h_i) - \mu_{\log})^2$$

$$\text{median} = \exp\left(\mu_{\log}\right)$$

$$\text{mode} = \exp\left(\mu_{\log} - V_{\log}\right)$$

Thickness statistics

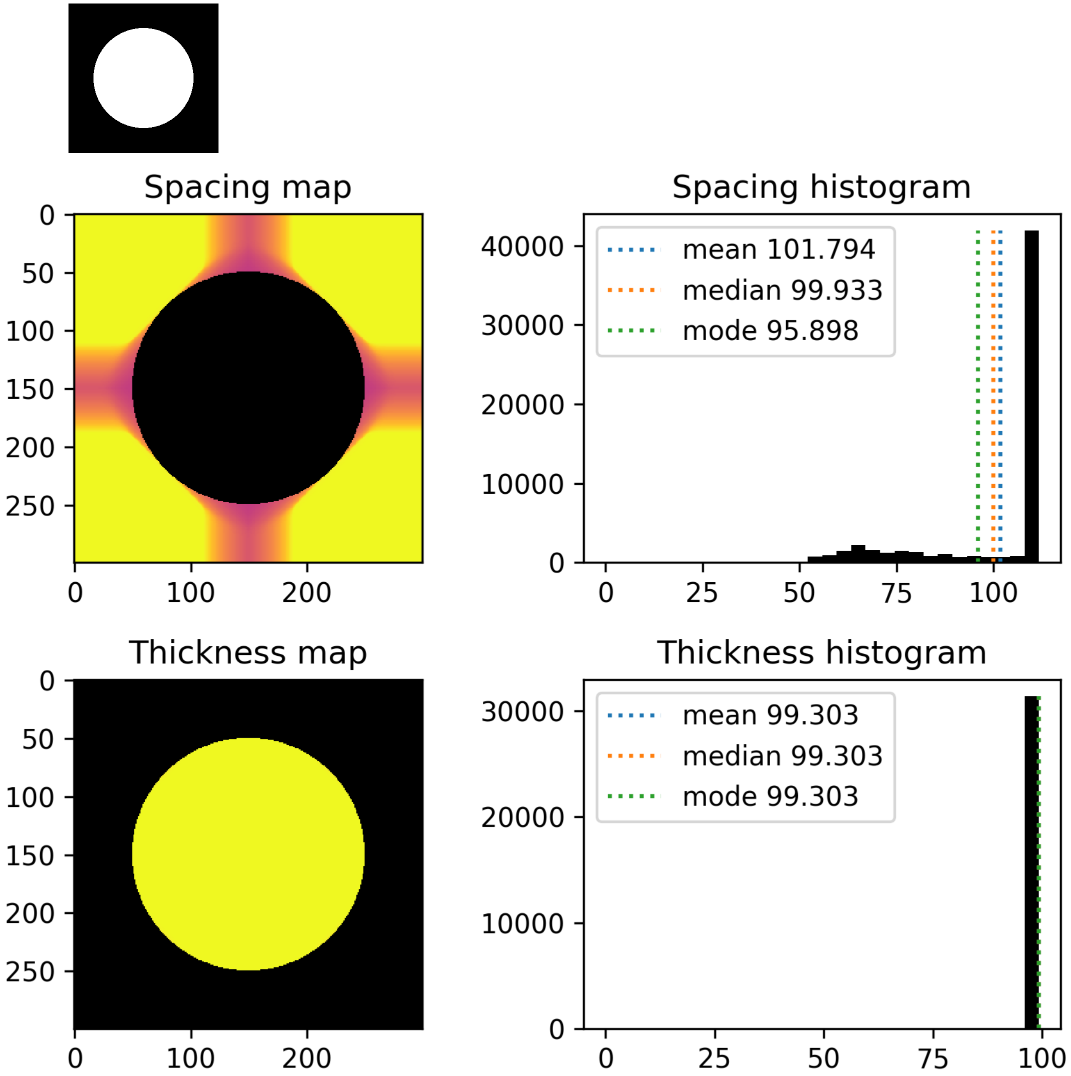

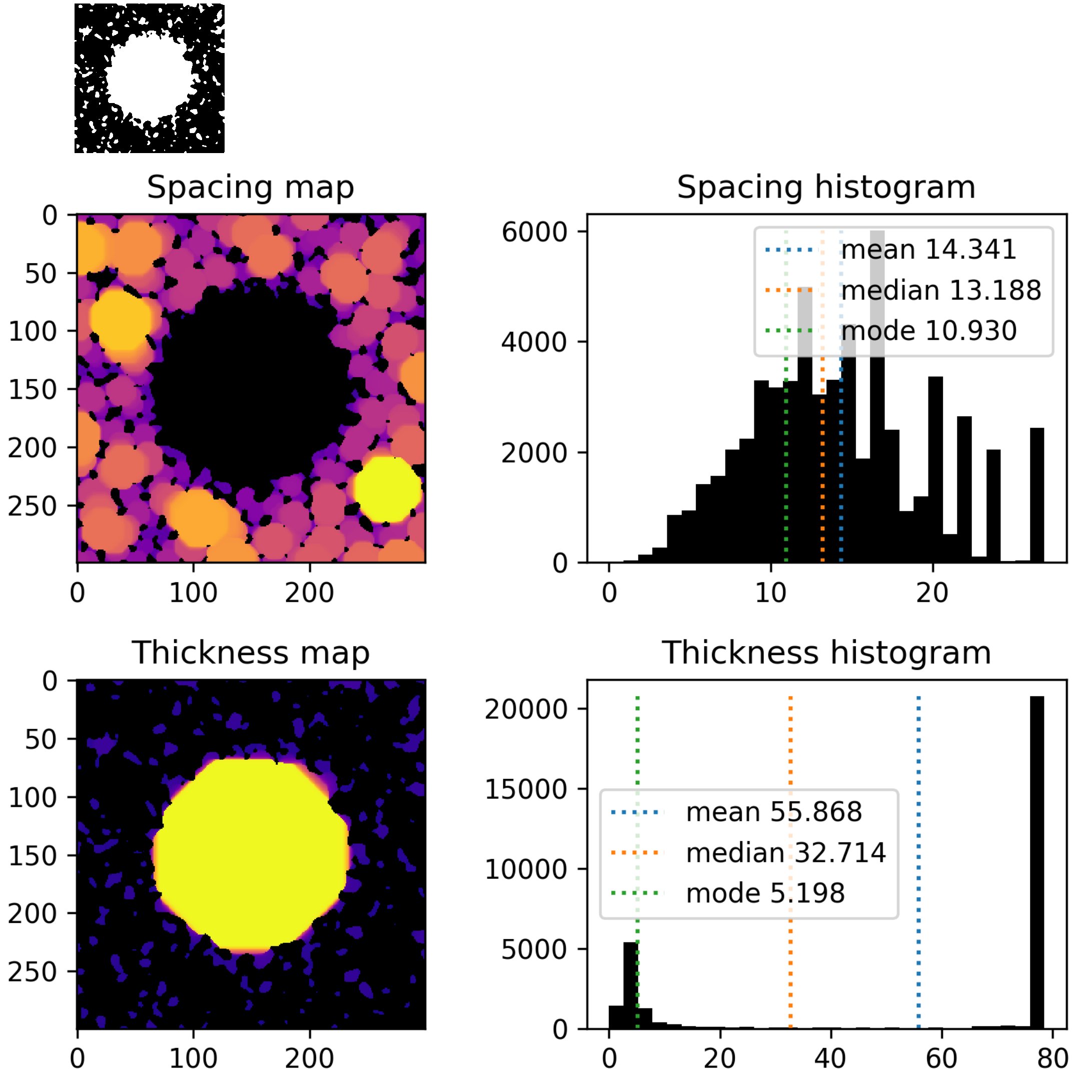

Quantitative measures

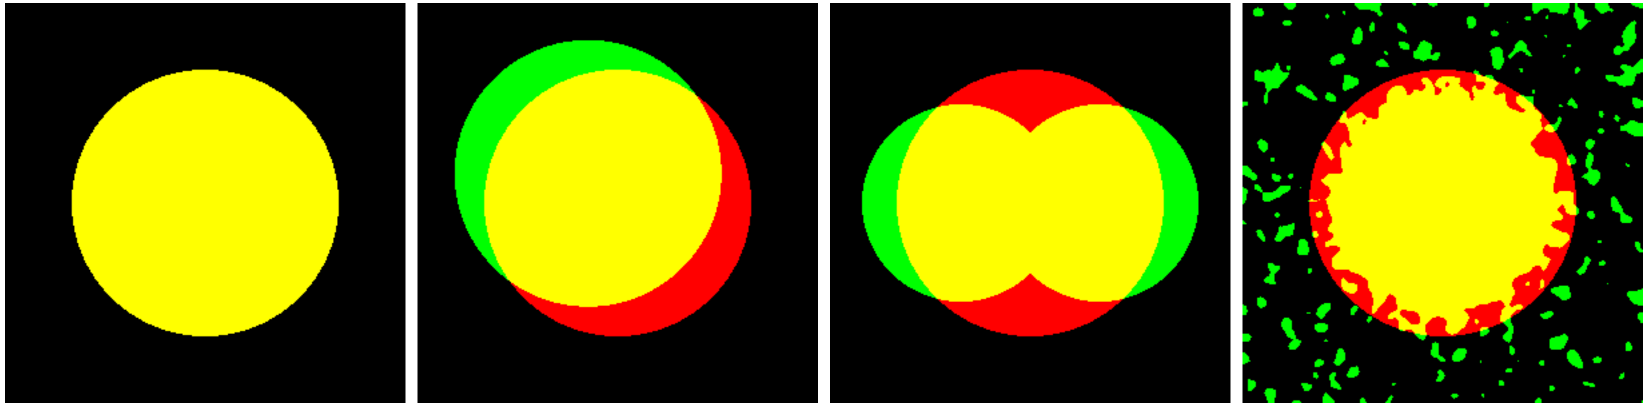

Boundary length ratio	1	1	1.11	6.75
No. of comp. – foreground	1	1	1	135
No. of comp. – background	1	1	1	4
Median thickness	97.8	98.9	72.3	31.4
Median spacing	100.0	105.8	72.1	13.0

Segmentation of kidney data

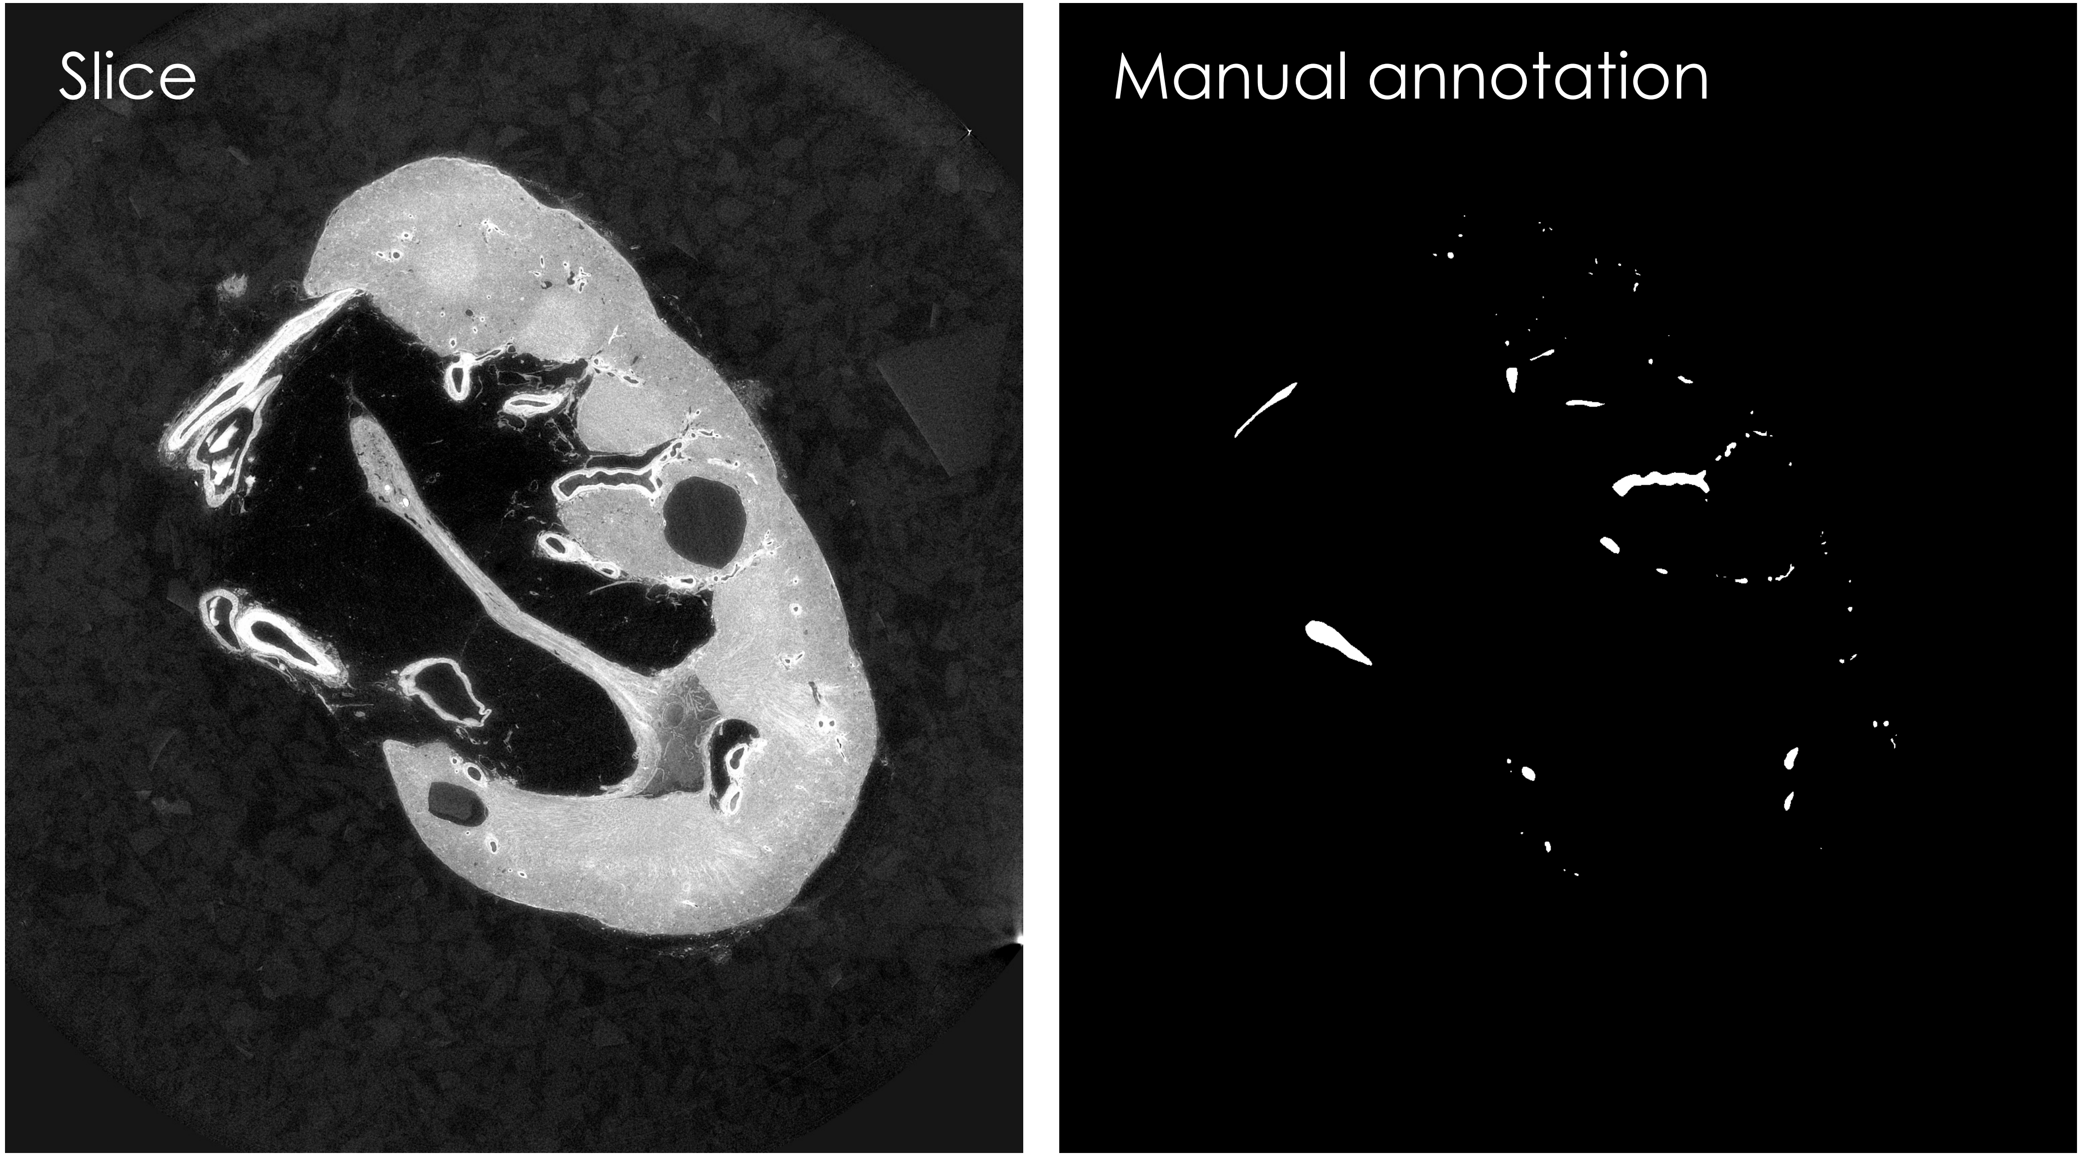

Segmentation of kidney data – U-Net

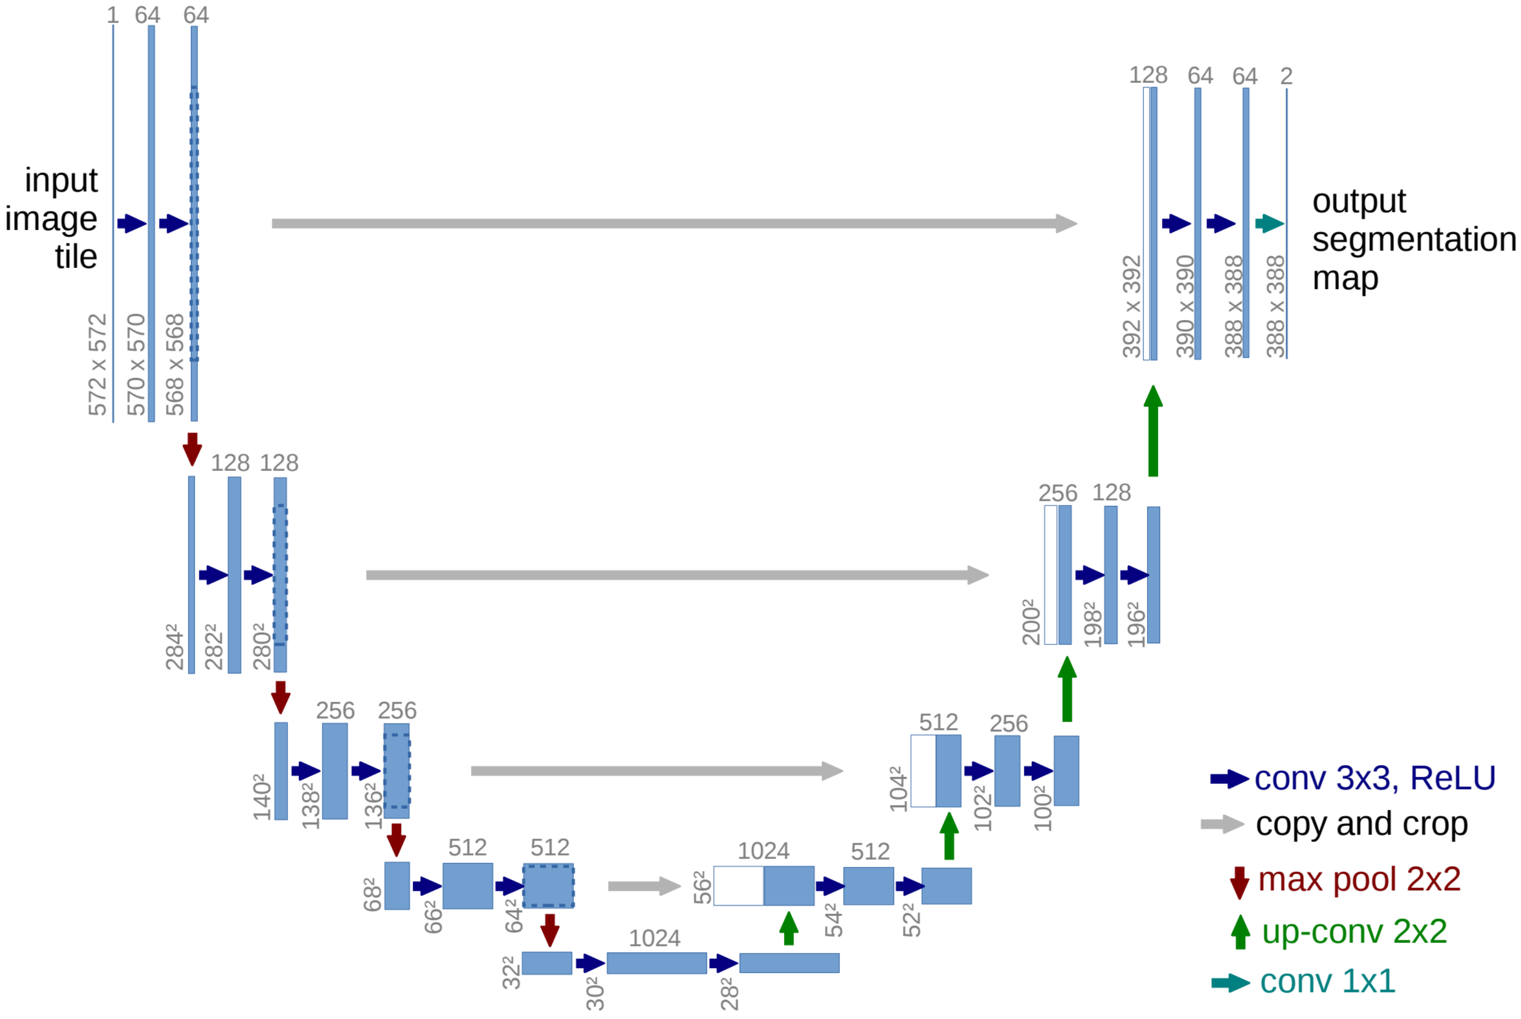

Kidney example

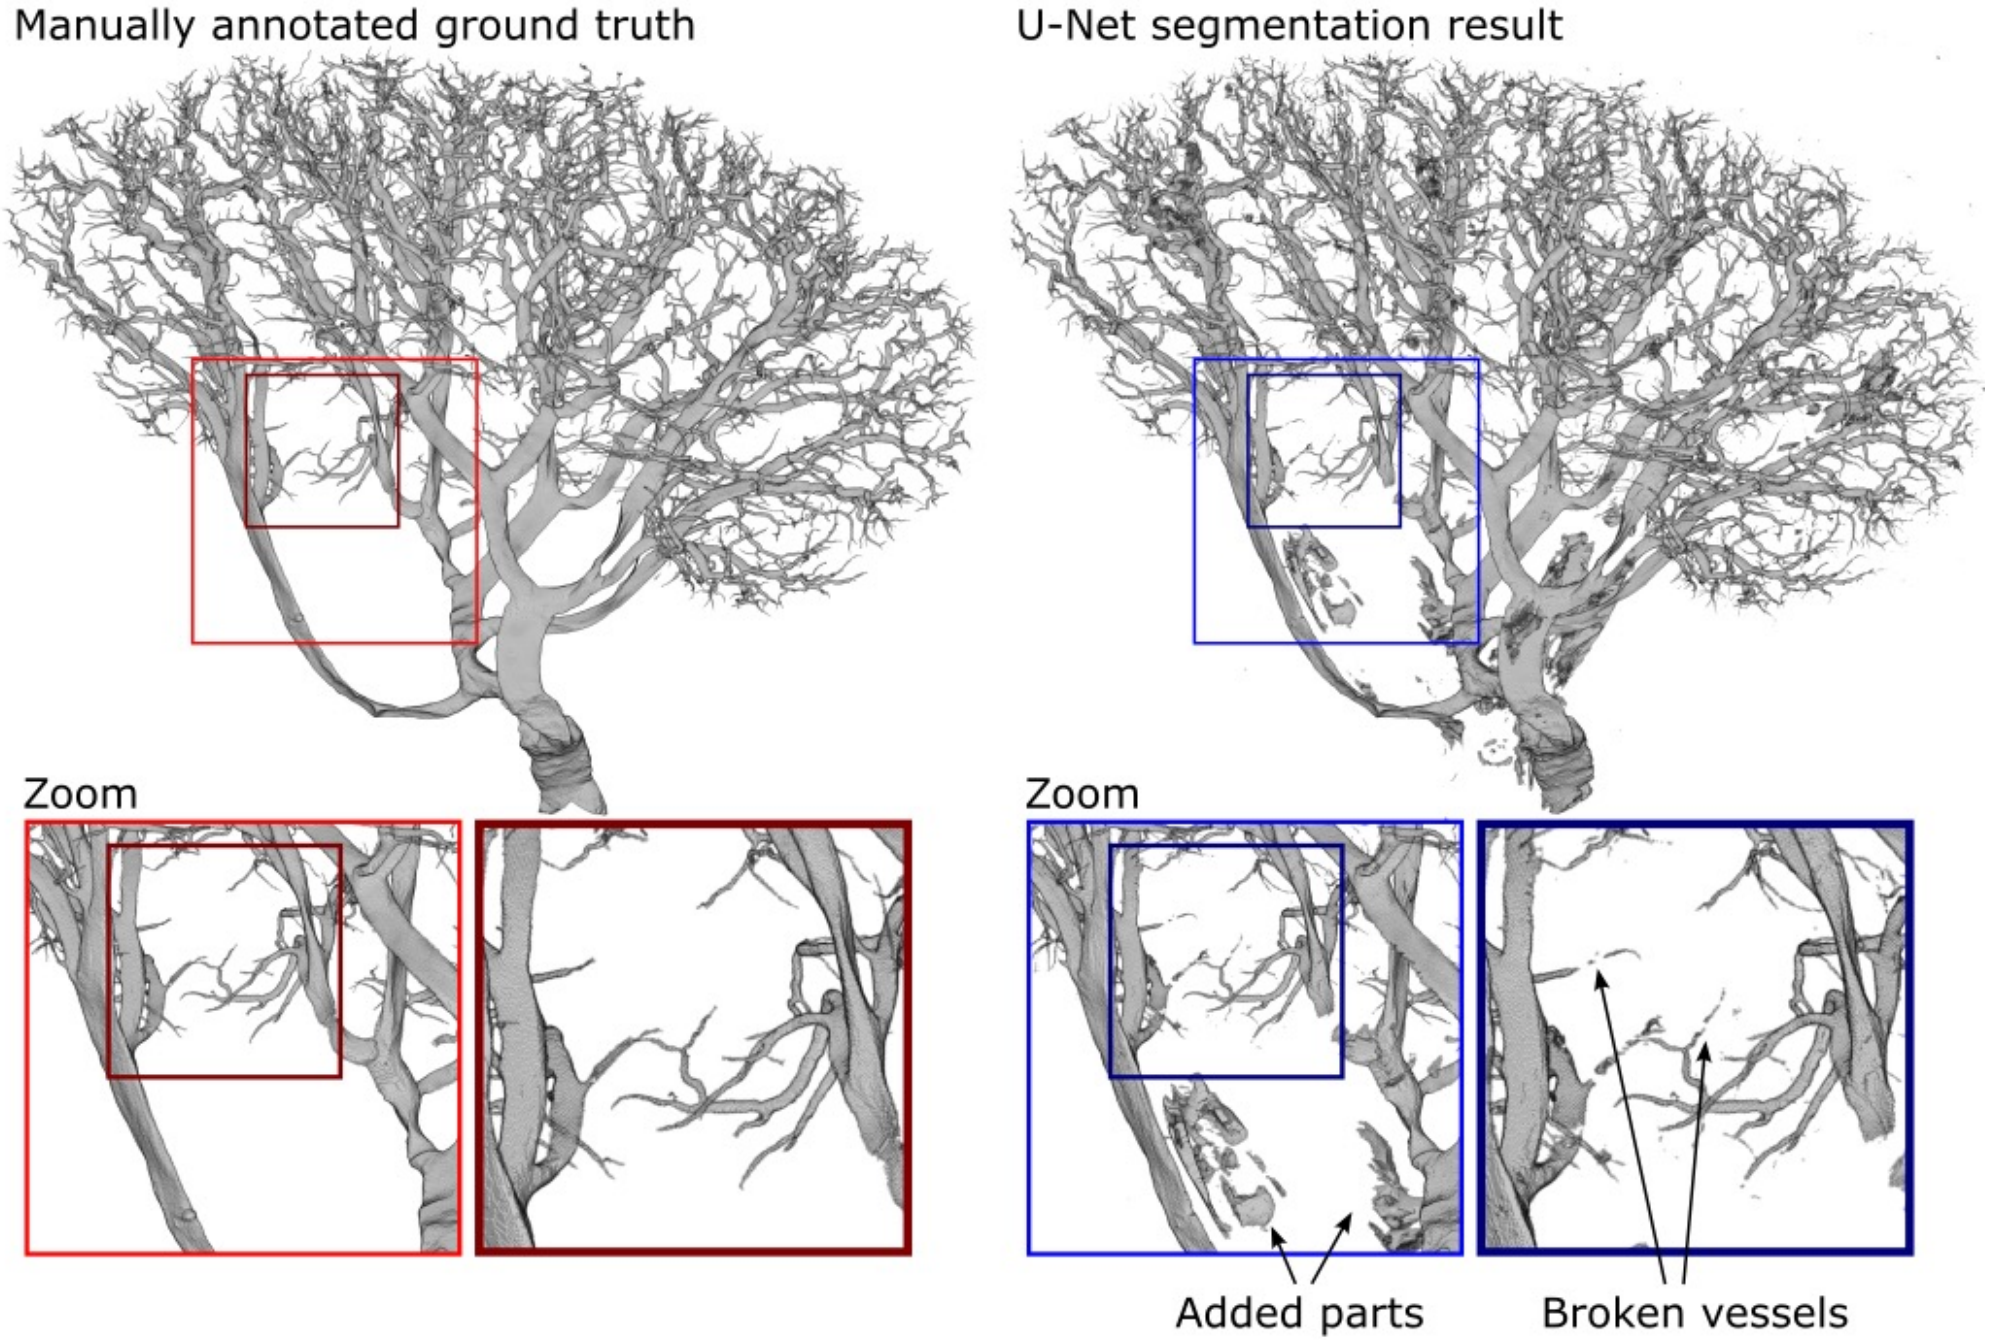

Per slice stats:

Dice:	0.740
# comp. ratio:	1.21
Thick. ratio:	0.919
Space. ratio:	0.970

Conclusion (part one)

- Segmentation measures may not reflect segmentation quality (including accuracy, Dice, and IoU)
- Visualizing and inspecting results is essential
- Need for additional research
  - Quantitative analysis methods
  - Segmentation techniques
  - Etc.

How do we get people to work for free?

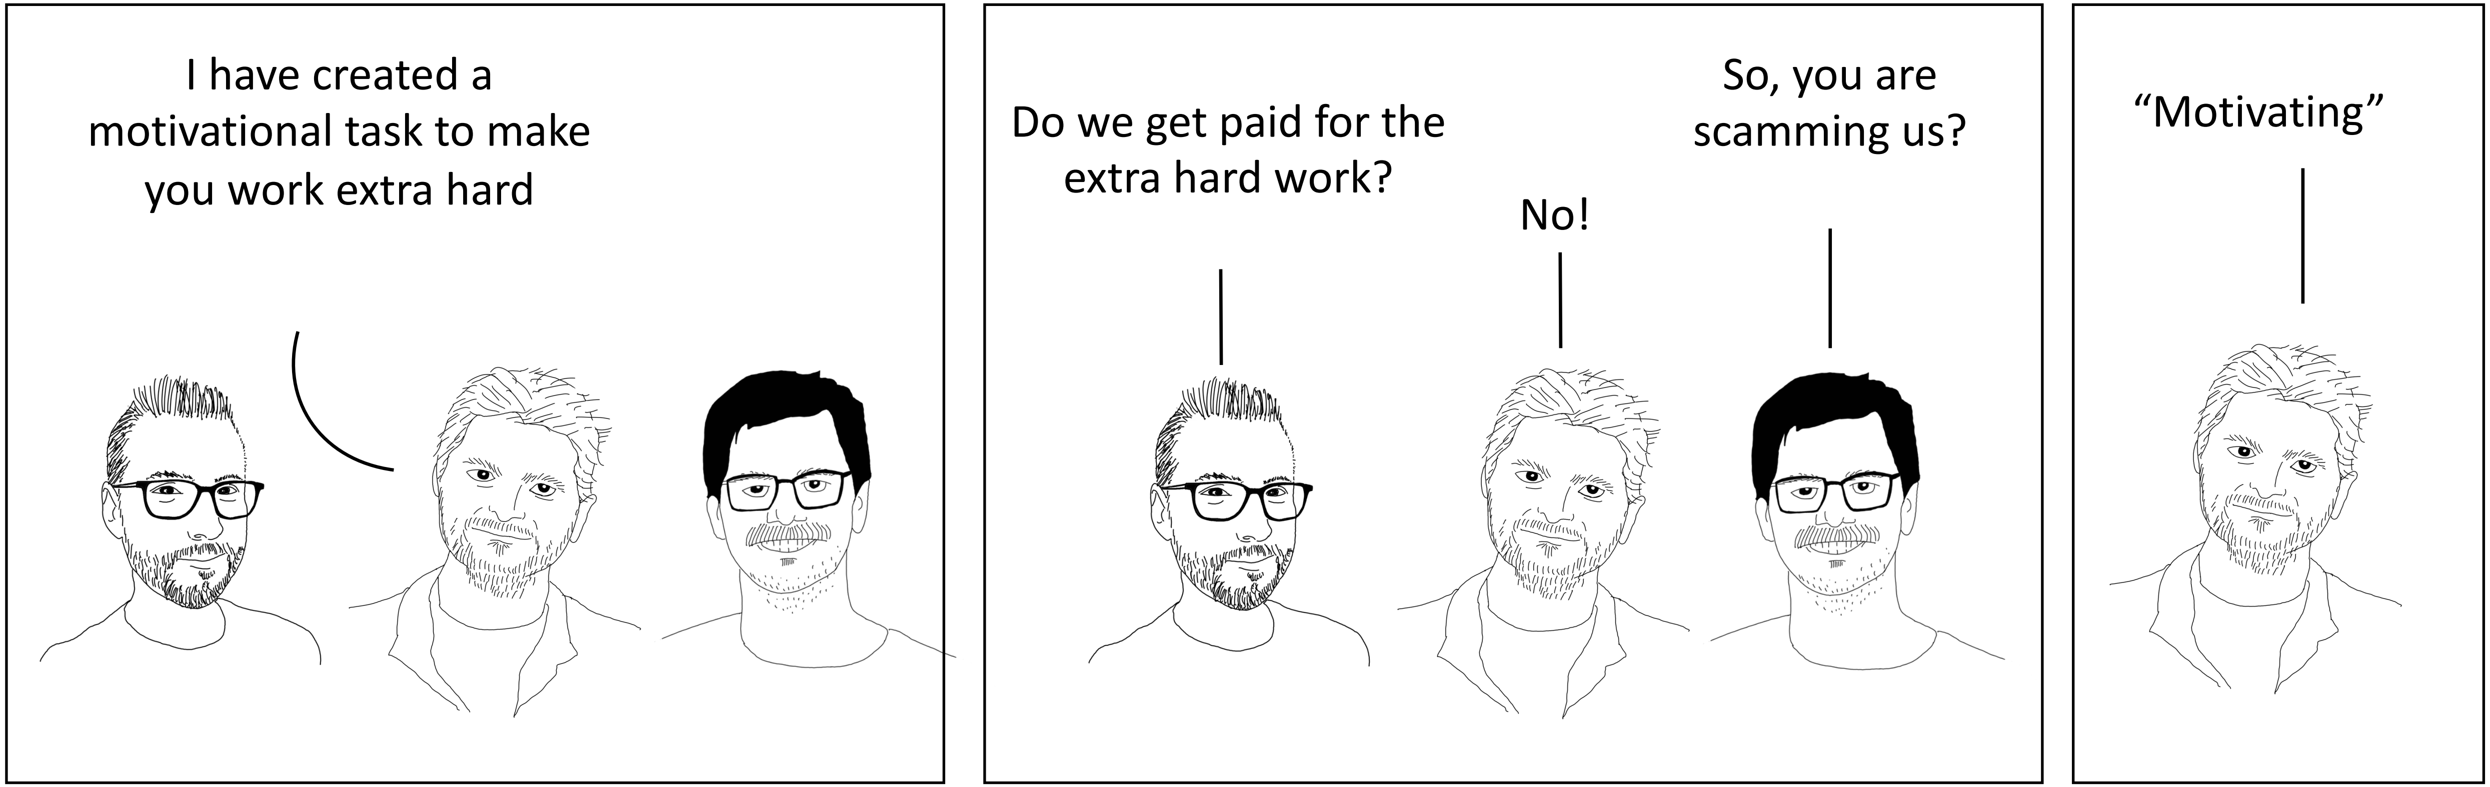

Modified from Dilbert

Idea

- Create a dataset – a benchmark for deep learning
  - Inspired by MNIST and CIFAR
  - Volumetric images
  - Easy to use
  - Large scale
- Get other researchers to work on the data
- Collect results and use for other types of problems
- I got the idea, that bugs would be ideal
  - Collaboration with the Natural History Museum
  - Bugs are complex in shape
  - People know about bugs and can relate to them
  - No GDPR

Data collection

Lab

Our freezer at home

Development of packaging


VIGTIGT (til modtager)
Er pakken sendt med post ved
en temperatur på 5 grader
eller under,skal den i
stuetemperatur, 2 timer
før UDPAKNING


BEKO
simplestore
2.0

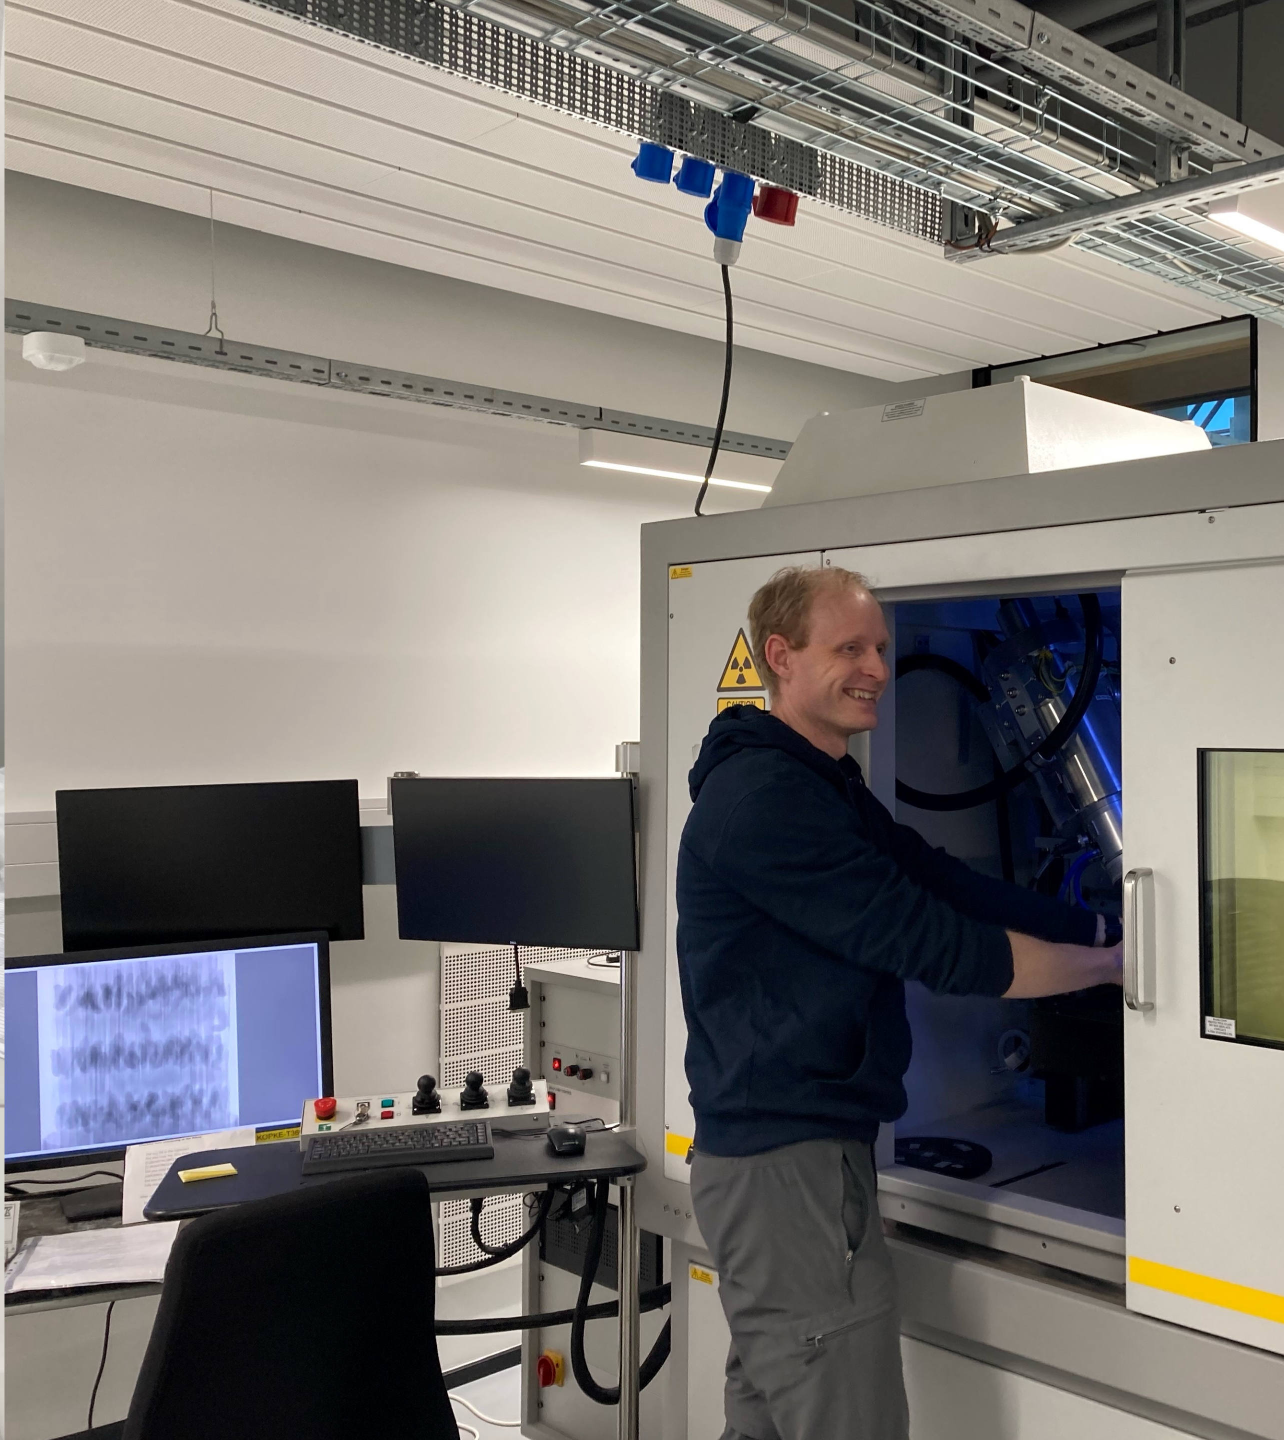
CAUTION

Initial data

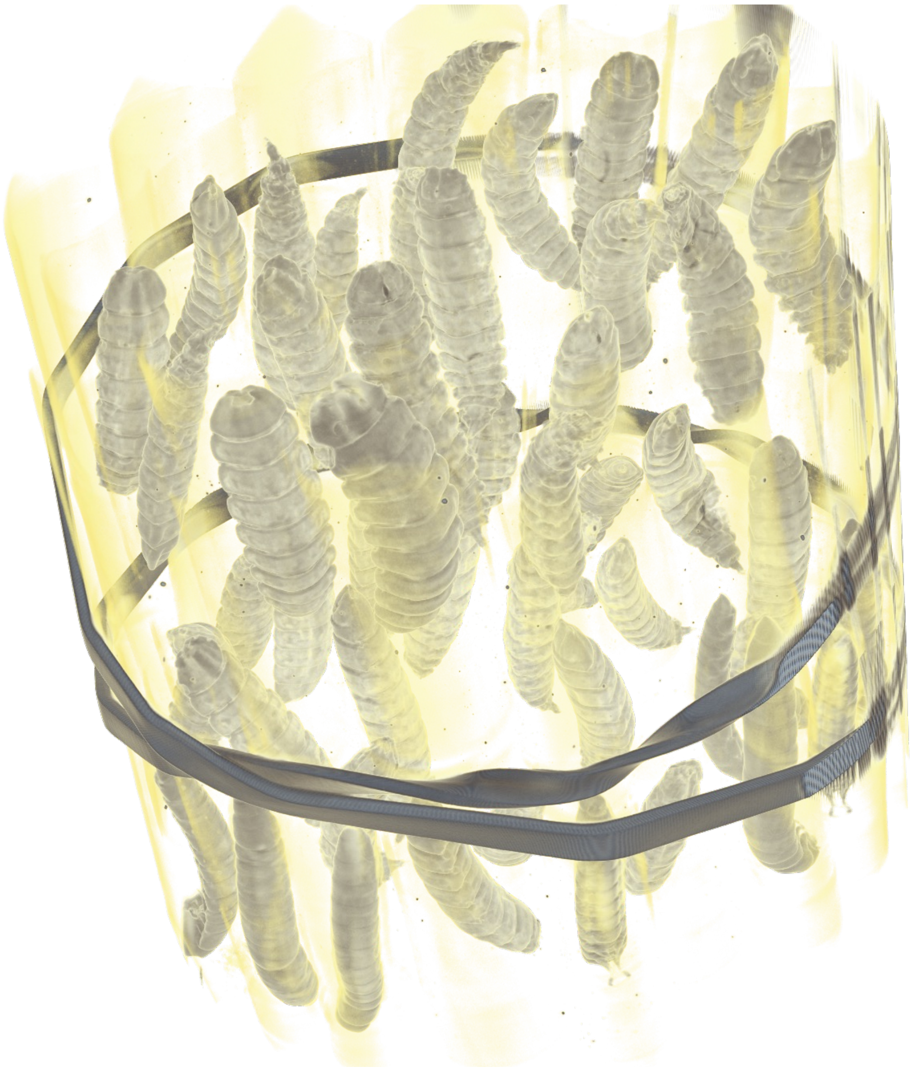

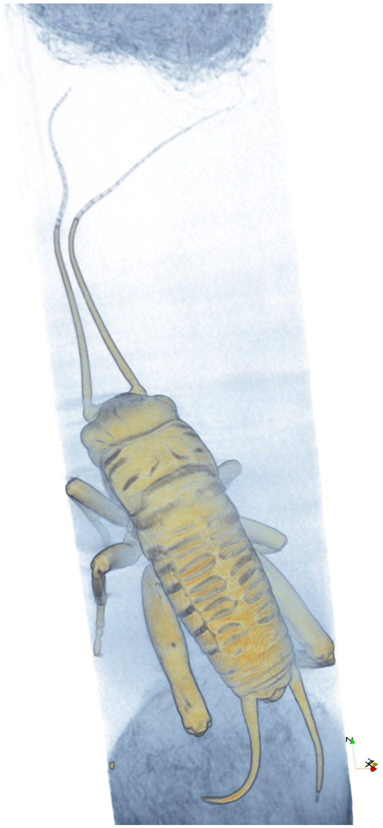

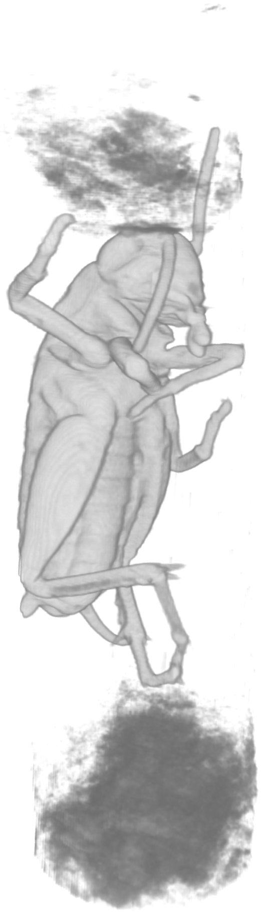

Dataset for classification

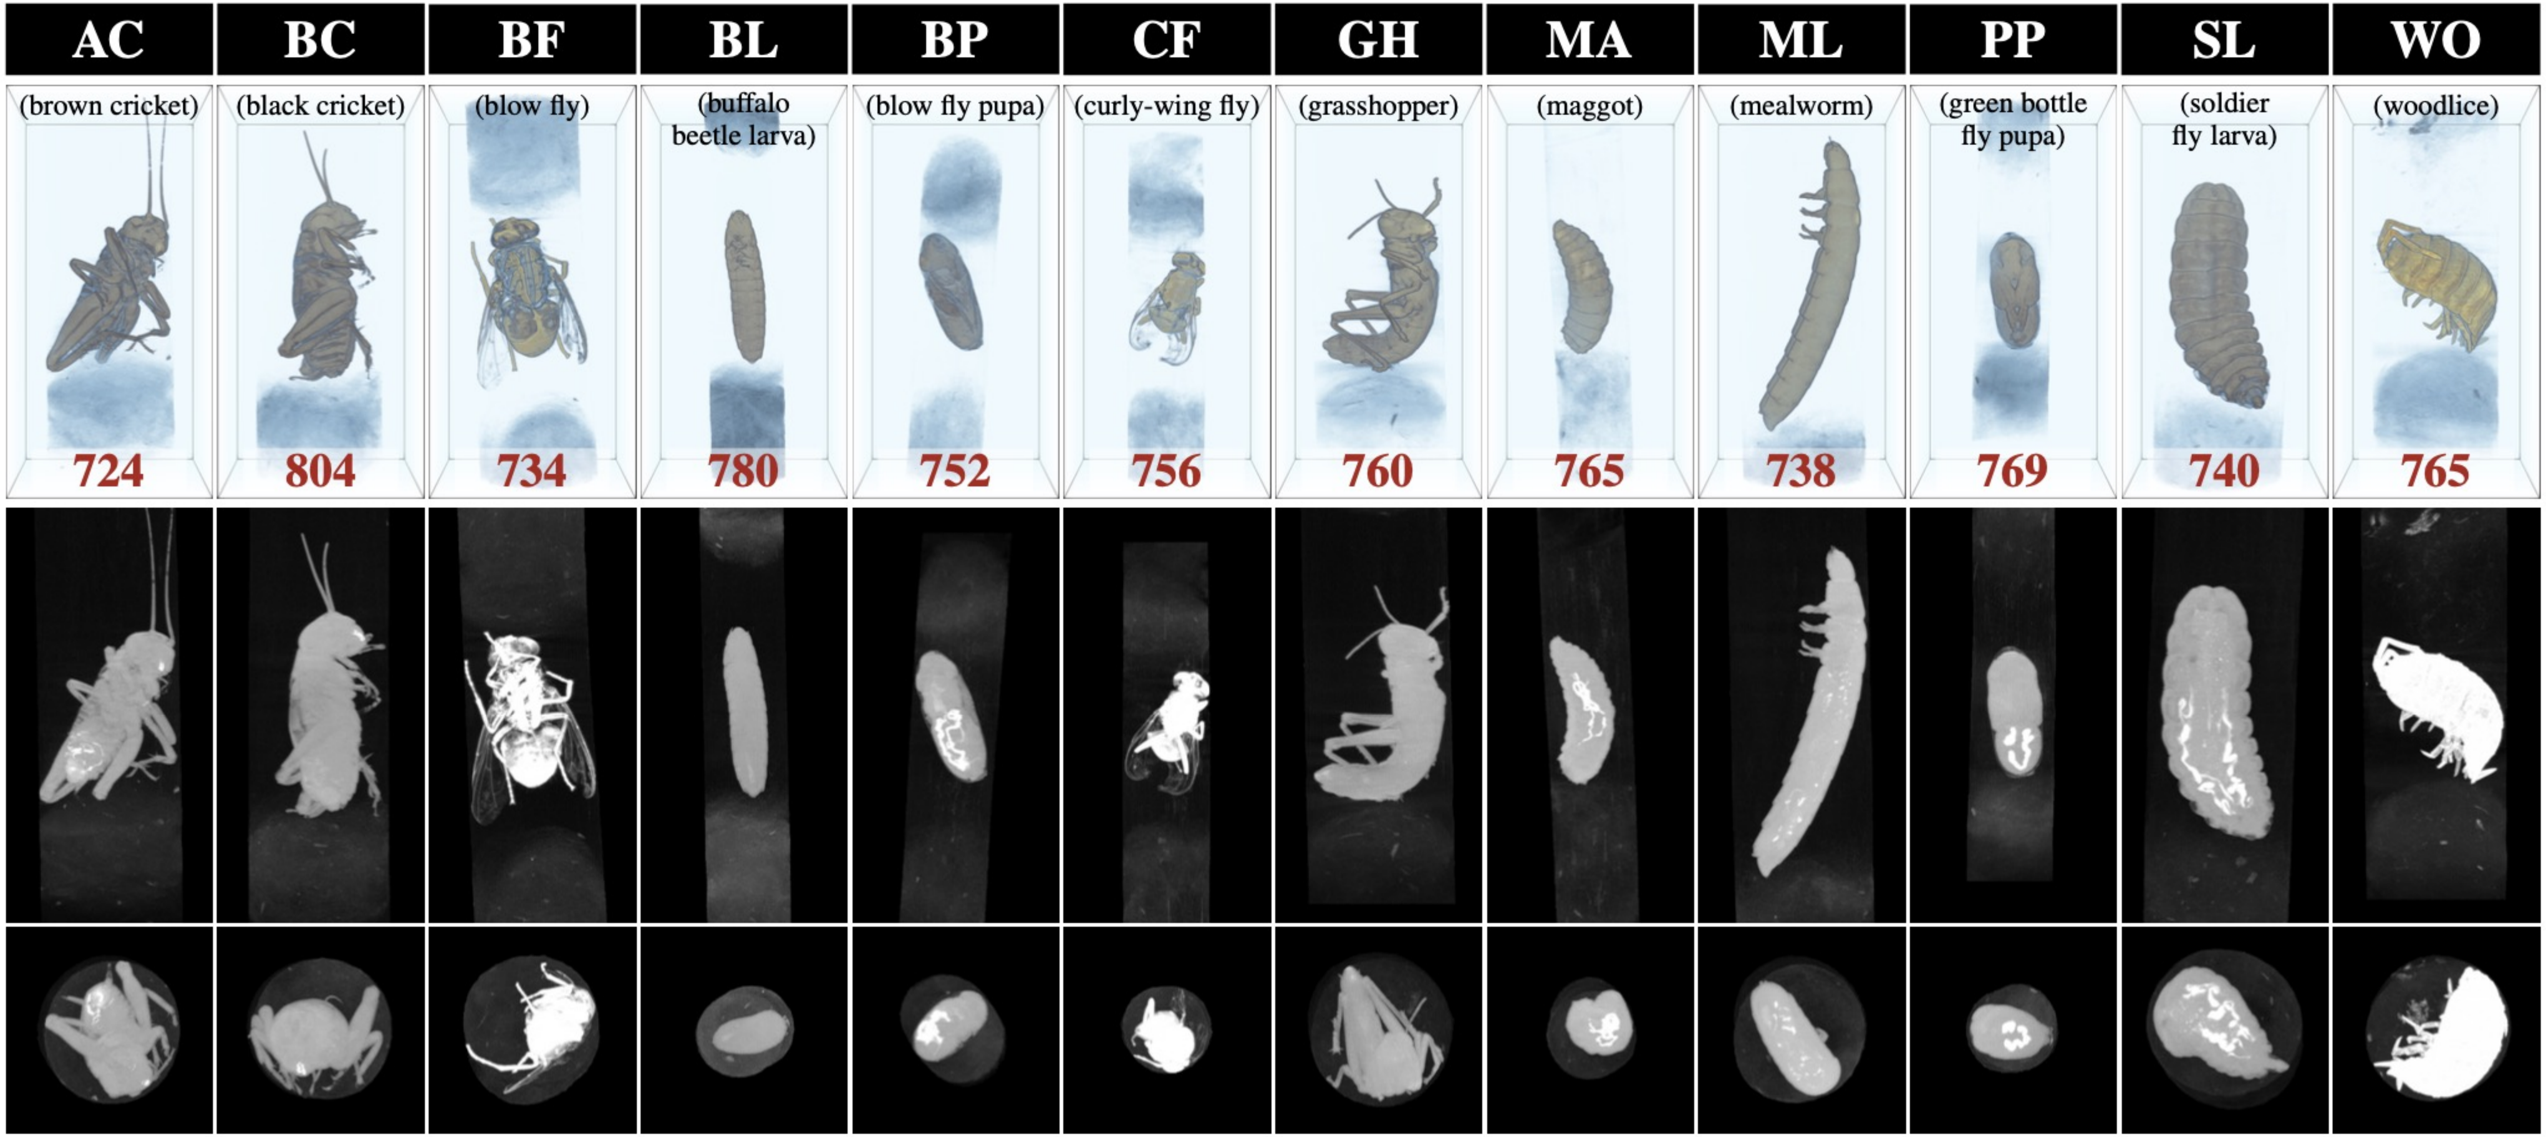

Results - classification

- Almost 100% correct
- Small differences between crickets
- Not much to talk about – solved problem!
- What to do…?

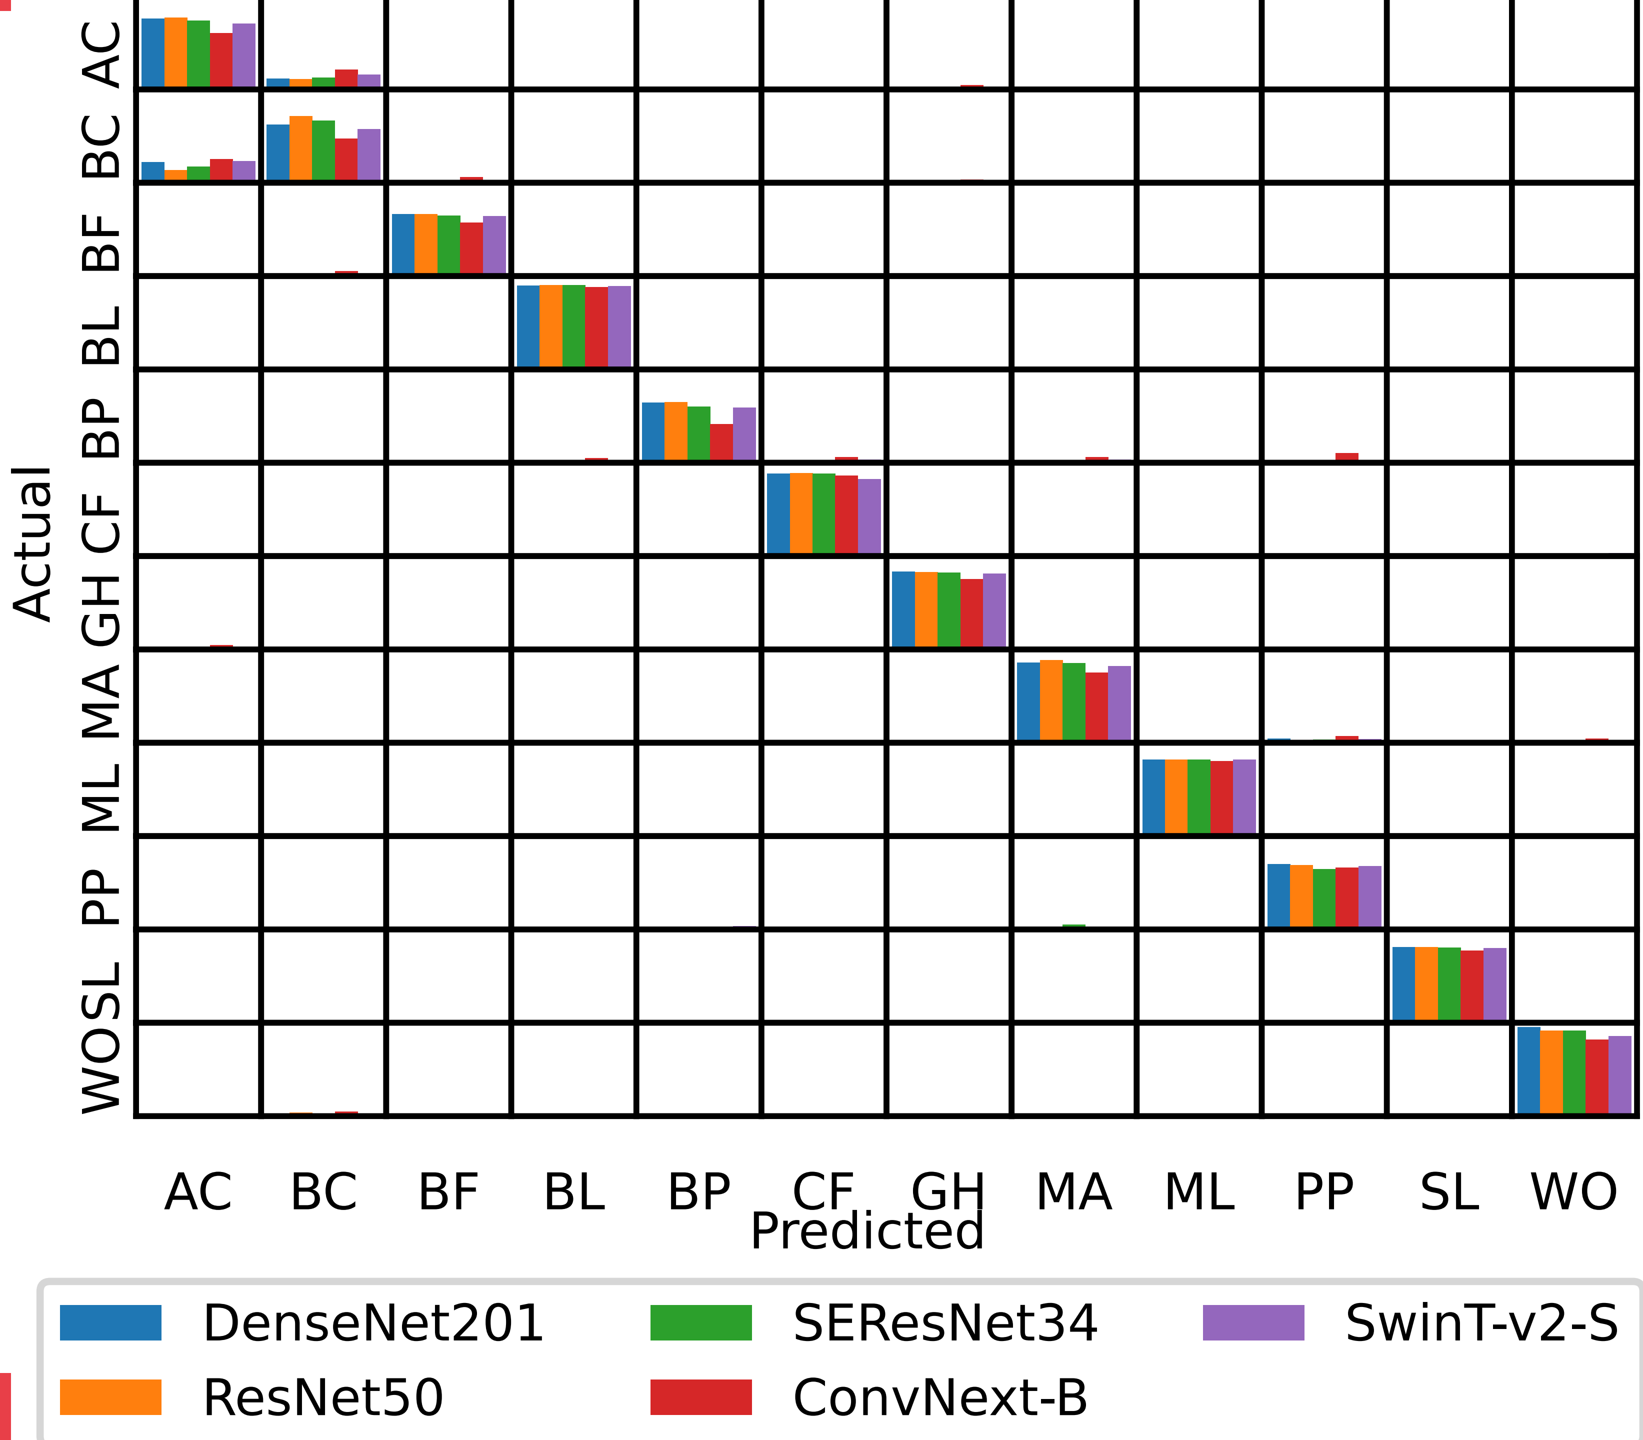

Create a more difficult problem

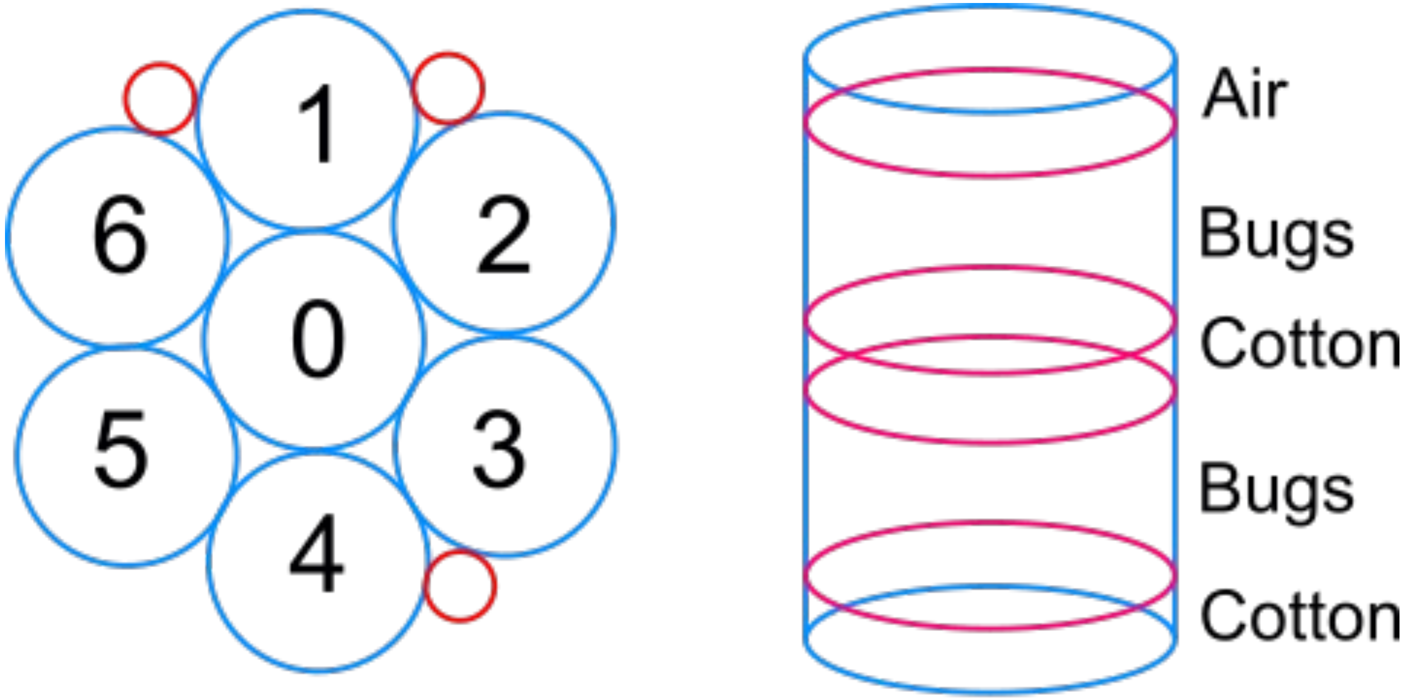

Bundle of test tubes

Single test tubes

Small Change – Catastrophic Consequences

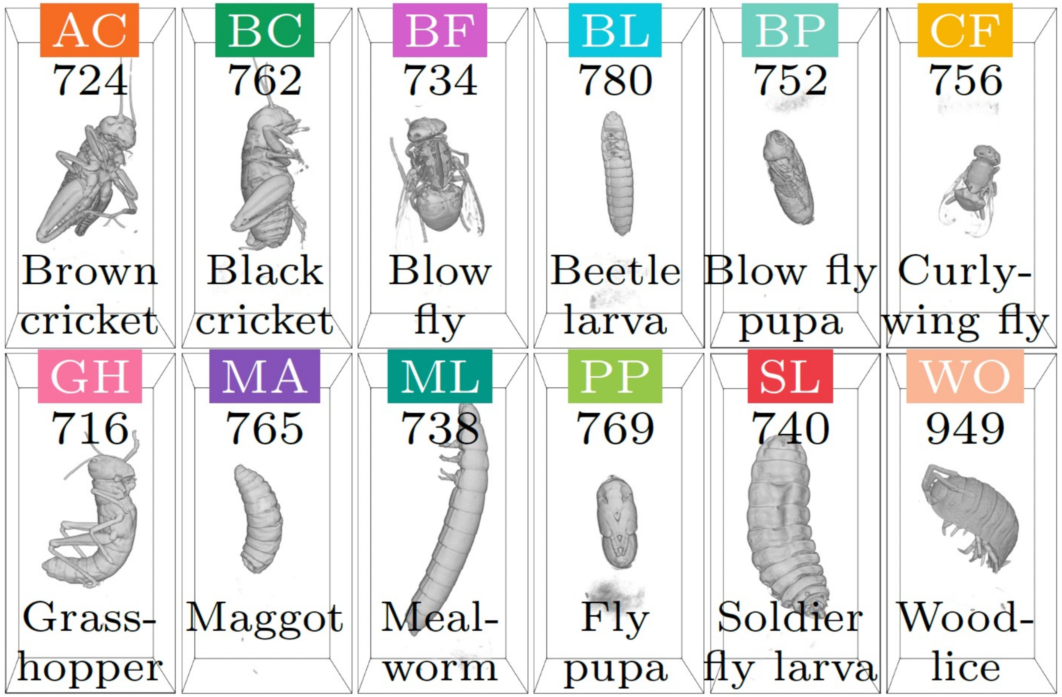

Training
Single objects (µCT volumes)

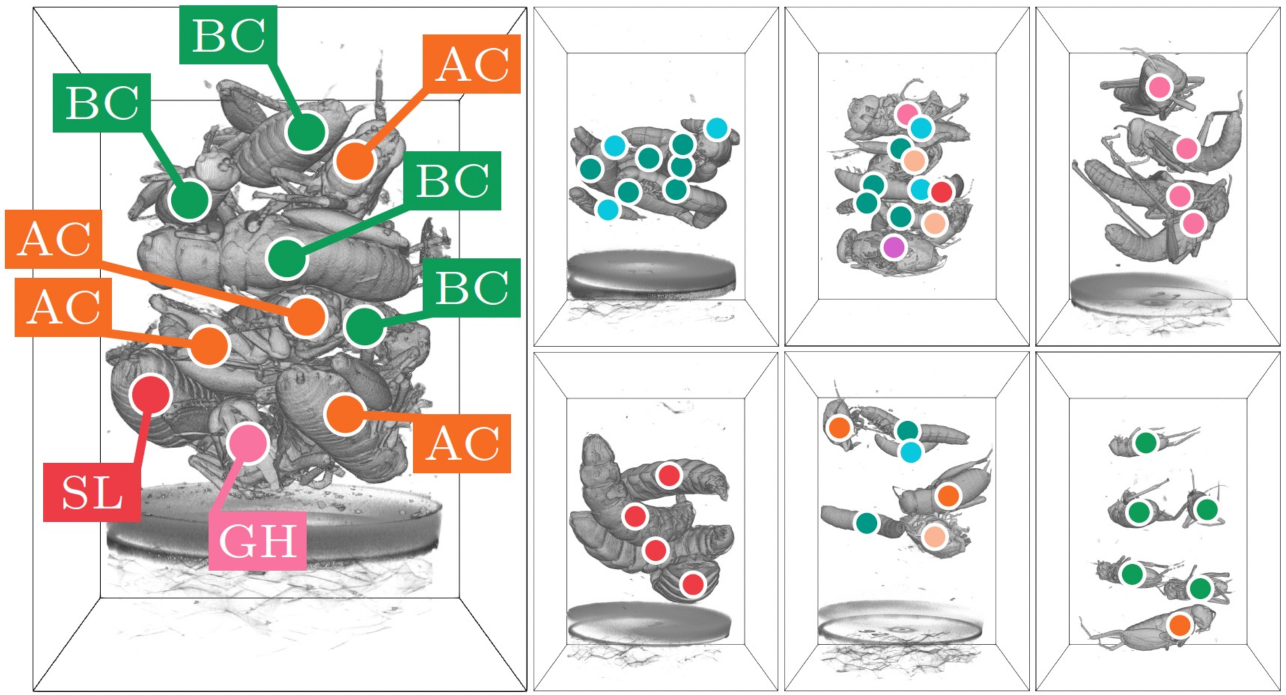

Test
Mixed objects (µCT volumes)

Source: Jensen et al., in prep.

Data

Individual – simple to annotate

Mixes – complex to annotate

Deep learning benchmark data

- Topic: Object detection and segmentation
- Problem: Annotation is difficult and time-consuming
- Idea: Create a setup where annotation is easy
- Dataset for deep learning:
  - 9154 individual bugs – reflecting shape variation
  - 388 volumes of mixtures
  - Scans with background material (no bug scans)
  - Publish data, baseline solution, and Kaggle Challenge

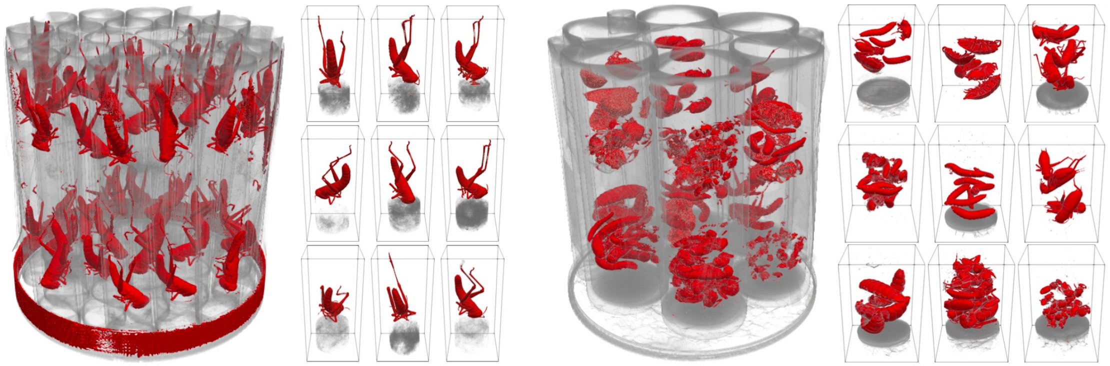

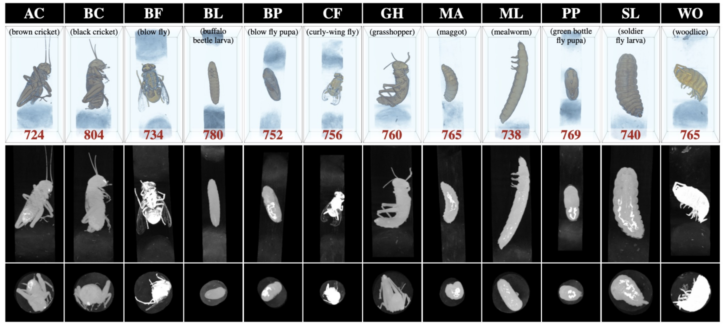

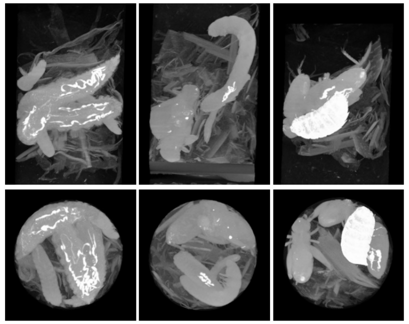

Results on BugNIST challenge

		Without class info.			With class info.		
Model	- Trained on	F1-Score	Precision	Recall	F1-Score	Precision	Recall
U-Net [8]	- Single bugs	0.59±0.18	0.68±0.17	0.60±0.25	0.11±0.16	0.13±0.19	0.11±0.16
	- Synth. mixes	0.60±0.12	0.52±0.17	0.82±0.19	0.11±0.08	0.09±0.08	0.14±0.10
	- Crowded s.m.	0.54±0.13	0.39±0.14	0.97±0.07	0.10±0.08	0.07±0.06	0.18±0.14
Faster R-CNN [54]	- Single bugs	0.56±0.20	1.00±0.04	0.42±0.21	0.11±0.10	0.21±0.19	0.08±0.08
	- Synth. mixes	0.68±0.16	0.63±0.21	0.83±0.22	0.16±0.14	0.15±0.14	0.19±0.18
	- Crowded s.m.	0.27±0.17	0.81±0.24	0.18±0.16	0.06±0.09	0.19±0.30	0.04±0.06
nn-Detection [4]	- Single bugs	0.03±0.02	0.02±0.01	0.91±0.15	0.00±0.01	0.00±0.00	0.09±0.16
	- Synth. mixes	0.50±0.14	0.41±0.17	0.80±0.22	0.10±0.06	0.08±0.06	0.16±0.11
	- Crowded s.m.	0.60±0.19	0.77±0.23	0.58±0.27	0.11±0.12	0.15±0.16	0.11±0.12

FGVC Kaggle challenge @ CVPR

CVPR SEATTLE, WA JUNE 17-21, 2024

BugNIST2024 Volumetric Out-of-Context Detection

Patrick M. Jensen, Vedrana A. Dahl, Rebecca Engberg, Carsten Gundlach, Hans Martin Kjer, Anders B. Dahl

BugNIST

Background: Labeling densely packed 3D objects for detection is time-consuming.

Suggestion: Train detection on isolated, individually scanned objects, where labels are easy to obtain.

Problem: Inferring on densely packaged objects introduces domain shift: out-of-context.

Example objects: Bugs. Due to size suitable for handling and scanning, complex shape, and availability.

Data collection: Bugs placed in tubes, individually or in mixtures (shown below).

Aim: Promote the development of new deep-learning-based methods for 3D volumetric imaging.

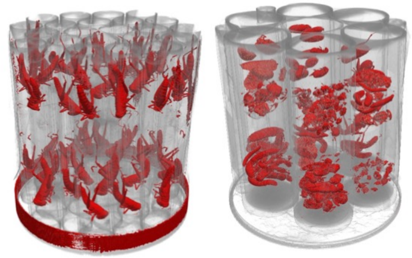

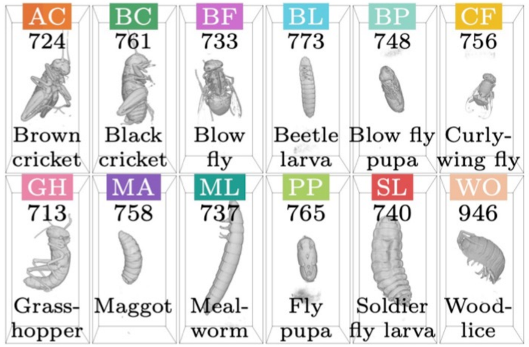

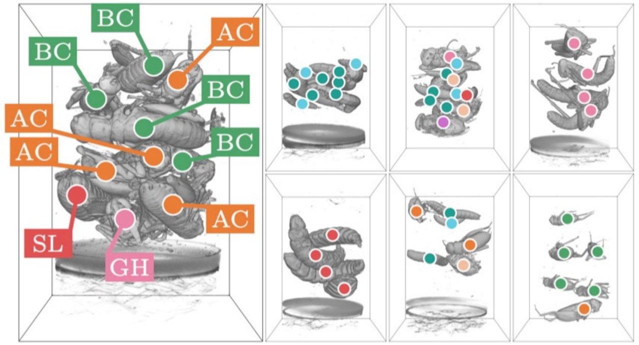

BugNIST dataset

- Individual bug volumes: 9185 micro-CT volumes with single bug, 12 bug classes
- Bug mixes: 388 volumes of densely packed bugs, each bug annotated with bug class and center point

Out-of-context detection

- Training: only individual bugs. (These may be used to create synthetic mixes.)
- Testing: bug mixes.

Bugs have the same appearance in the source (individual) and target (mixes) domain, but the surrounding context is different.

kaggle challenge

Participation

- 108 Entrants
- 31 Participants
- 16 Teams
- 91 Submissions

Prizes & Awards

- Kudos
- Joint paper

Team	Public leaderboard	Private leaderboard
Baseline	0.11102	0.12711
Winning team	**0.55318**	**0.54899**
Second place	0.45457	0.44450

Reflections: Limited overfitting – similar performance on private and public leaderboards.

- 11 teams beat the baseline
- 4 teams below the baseline

Approach (baseline)

- U-Net with a depth of three trained on individual bugs
- Detection – center points of connected components

Approach (winning team)

- U-Net with a depth of three trained on synthetic mixes
- On-the-fly synthesis and augmentation
- Synthetic mixes inspired by Tetris
- Random augmentations
- Cross-entropy and non-background Dice loss
- Detection – center points of connected components

Winning team: Pioneer Bugsifiers

Assistant Professor Mostafa Mehdipour Ghazi
PhD, ML Developer Torkan Gholamalizadeh
University of Copenhagen, AI Pioneer Center

Data avaliable

Webpage for the BugNIST dataset

View our competition on Kaggle
Kaggle/bugnist

Download **dataset** | Get arXiv **paper** | Code on **GitHub**

BugNIST - dataset for volumetric analysis

The BugNIST dataset is created to advance methods for classification and detection in 3D. It contains 9542 volumes where 9154 are of individual bugs and 388 are mixtures of bugs and other material. There are 12 types of bugs including larvae, pupae, insects, and woodlice.

In the BugNIST classification challenge, each volume containing a single bug must be classified as one of the 12 types. The original volumes are 900x450x450 voxels, and in addition, we provide the data at different resolutions by downscaling the original scans.

aims to benchmark classification and detection methods, and we have designed the detection challenge such that detection models are trained on scans of individual bugs and tested on bug mixtures. Models capable of solving this task will be independent of the context, i.e., the surrounding material. In cases where the context is unknown or changing, this is a great advantage, which is commonly occurring in 3D µCT.

What is BugNIST?

BugNIST is a volumetric dataset for object detection and segmentation. BugNIST has several features:

- Object detection in volumetric data
- Segmentation in volumetric data
- 3D µCT scanning of 12 bug classes
- 9154 volumes of individual bugs
- 388 volumes of bug mixtures with center point annotations
- Volume sizes: 900x450x450 (individual) and 900x650x650 (mixtures)
- Data available in sizes:
    - Original: 900x450x450 (individual) and 900x650x650 (mixtures)
    - Large: 512x256x256 (individual) and 512x370x370 (mixtures)
    - Medium: 256x128x128 (individual) and 256x185x185 (mixtures)
    - Small: 128x64x64 (individual) and 128x92x92 (mixtures)
    - Tiny: 64x32x32 (individual) and 64x46x46 (mixtures)

People:

- Anders Bjorholm Dahl, DTU Compute
- Patrick Møller Jensen, DTU Compute
- Vedrana Andersen Dahl, DTU Compute
- Carsten Gundlach, DTU Physics
- Rebecca Engberg, DTU Compute
- Hans Martin Kjer, DTU Compute

Sponsors: Novo Nordisk Foundation, Villum Foundation

Page is maintained by abdahl and QIM team.

archive.compute.dtu.dk

LOG IN

zips

// public / projects / BugNIST3D / zips

Browse | Commits | Download .zip | Show Deleted Files

Name	Modified	Size
..		
bugnist_064.zip	about a month ago	86.68 MB
bugnist_128.zip	about a month ago	540.03 MB
bugnist_256.zip	about a month ago	4.29 GB
bugnist_512.zip	about a month ago	37.29 GB
bugnist_900.zip	about a month ago	237.65 GB
mixed_064.zip	about a month ago	15.62 MB
mixed_128.zip	about a month ago	125.68 MB
mixed_256.zip	about a month ago	1.07 GB
mixed_512.zip	about a month ago	9.12 GB
mixed_900.zip	about a month ago	49.56 GB

https://abdahl.github.io/bugnist/ https://archive.compute.dtu.dk/files/public/projects/BugNIST3D/zips

Publication

arXiv v1:

- **Title:** BugNIST -- A New Large Scale Volumetric 3D Image Dataset for Classification and Detection
- **Authors:** Anders Bjorholm Dahl, Patrick Møller Jensen, Carsten Gundlach, Rebecca Engberg, Hans Martin Kjer, Vedrana Andersen Dahl
- **Rejected** at ICCV 2023

arXiv v2:

- **Title:** BugNIST -- a Large Volumetric Dataset for Object Detection under Domain Shift
- **Authors:** Patrick Møller Jensen, Vedrana Andersen Dahl, Carsten Gundlach, Rebecca Engberg, Hans Martin Kjer, Anders Bjorholm Dahl
- **Rejected** at CVPR 2024

arXiv v3:

- **Title:** BugNIST -- a Large Volumetric Dataset for Object Detection under Domain Shift
- **Authors:** Patrick Møller Jensen, Vedrana Andersen Dahl, Carsten Gundlach, Rebecca Engberg, Hans Martin Kjer, Anders Bjorholm Dahl
- **Accepted** at ECCV 2024!

BugNIST – a Large Volumetric Dataset for Object Detection under Domain Shift

Patrick Møller Jensen, Vedrana Andersen Dahl, Rebecca Engberg, Carsten Gundlach, Hans Marin Kjer, and Anders Bjorholm Dahl

Technical University of Denmark, Kgs. Lyngby, Denmark
{patmjen,vand,reen}@dtu.dk,cagu@fysik.dtu.dk,{hmjk,abda}@dtu.dk

Abstract. Domain shift significantly influences the performance of deep learning algorithms, particularly for object detection within volumetric 3D images. Annotated training data is essential for deep learning-based object detection. However, annotating densely packed objects is time-consuming and costly. Instead, we suggest training models on individually scanned objects, causing a domain shift between training and detection data. To address this challenge, we introduce the BugNIST dataset, comprising 9154 micro-CT volumes of 12 bug types and 388 volumes of tightly packed bug mixtures. This dataset is characterized by having objects with the same appearance in the source and target domains, which is uncommon for other benchmark datasets for domain shift. During training, individual bug volumes labeled by class are utilized, while testing employs mixtures with center point annotations and bug type labels. Together with the dataset, we provide a baseline detection analysis, with the aim of advancing the field of 3D object detection methods.

Keywords: Volumetric Dataset, Benchmark, Volumetric Object Detection, Domain Shift.

1 Introduction

Our work on domain shift in volumetric 3D images is motivated by the need for labeled data to train supervised deep learning models for volumetric imaging. The problem that we aim to solve is object detection and classification. We propose to label images of objects scanned as isolated entities as the source domain and use these as a basis to train models for object detection and classification in a complex context of mixed objects and other materials as the target domain. The effort needed for obtaining labels in the two domains is significantly different. If objects are isolated, they can be automatically labeled whereas mixed objects require expensive manual labeling. Automating the labeling based on isolated objects, however, leads to a domain shift between the data in the source domain for training and the data in the target domain for detection and classification. This domain shift is special because the appearance of the objects is the same

2 P.M. Jensen et al.

Future work

- Classification and object detection – not segmentation!
- Problem: Ground truth labels
  - Training data: Individual bugs that are automatically labeled
  - Test data: Mixes where annotation is manual and difficult

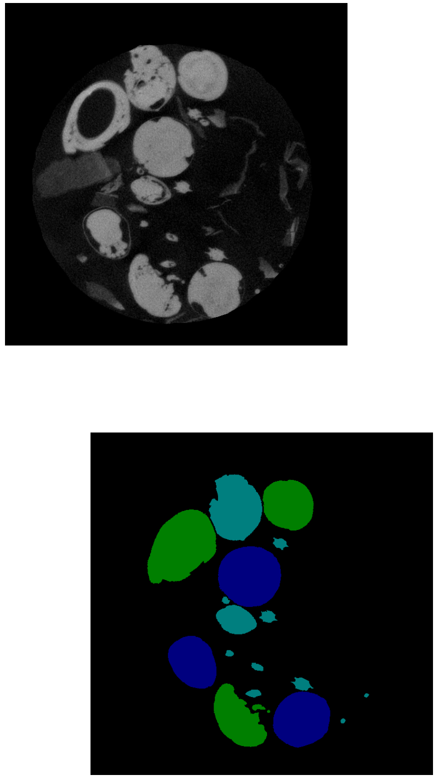

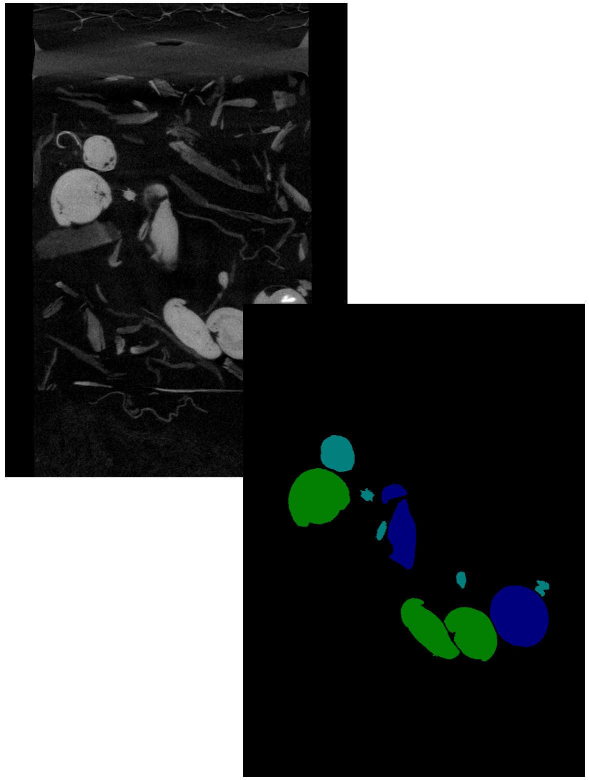

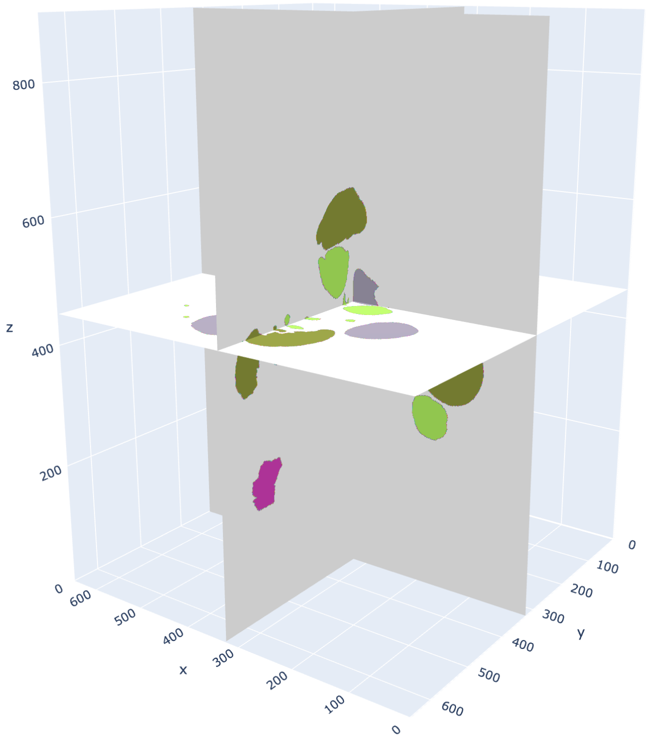

Conclusion (part two)

- Domain shift between individuals and mixes was surprisingly large
  - Minimal difference in appearance between individuals and mixes, yet standard deep learning fails!
  - Deep learning-based detection utilizes context
  - Potential research in methods that ignore context (initial investigations)
  - Potential research in generating context (strategy we tested and winners of Kaggle)
- Collecting data
  - Time consuming, exhausting, fun, frustrating…
  - Spend much time testing methods on preliminary data
  - The community expects extremely large datasets
- The data curation is ongoing and will continue
- Datasets with extensive investigations can get published in high-impact venues!

Jon Sporring: Exploring biological shape analysis through topology, geometry and statistics

The lecture will include a hands-on session. On https://sporring.github.io/ you will find the slides and a zip-file:

- https://sporring.github.io/bia2024/talk.pdf including a useful literature list
- https://sporring.github.io/bia2024/spatstat_bia2024.zip which includes demoRpy2.py and which has installation instructions for R, R-packages, and python packages:

```
# 1. Install R, which to my experience works best directly from  https://cran.r-project.org/
# then start R and install some packes:
#  install.packages("spatstat")
#  install.packages("lazyeval")
#  install.packages("GET")
…
```

Save time during the lecture and install R etc. in advance.

Questions?

Contact:

abda@dtu.dk