Scaling Annotations, Datasets, and Algorithms for Medical Image Analysis

Zongwei Zhou, PhD

Postdoc, Department of Computer Science
Johns Hopkins University, Baltimore, MD
P: 1-(480)738-2575 | E: zzhou82@jh.edu

Radiologists hate annotation, but computer scientists love it.

CT

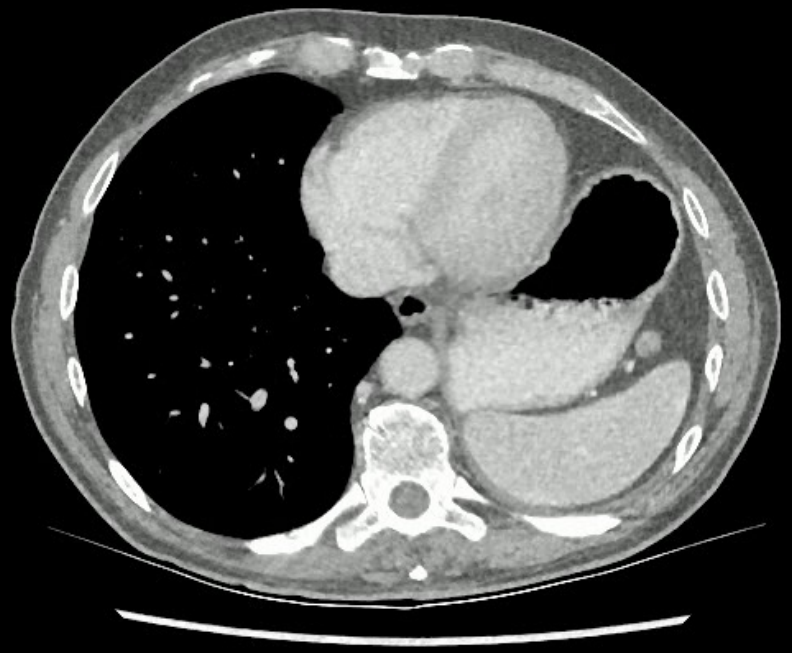

Ground Truth

annotated by human experts

Deep Learning

- Liver
- Liver tumor

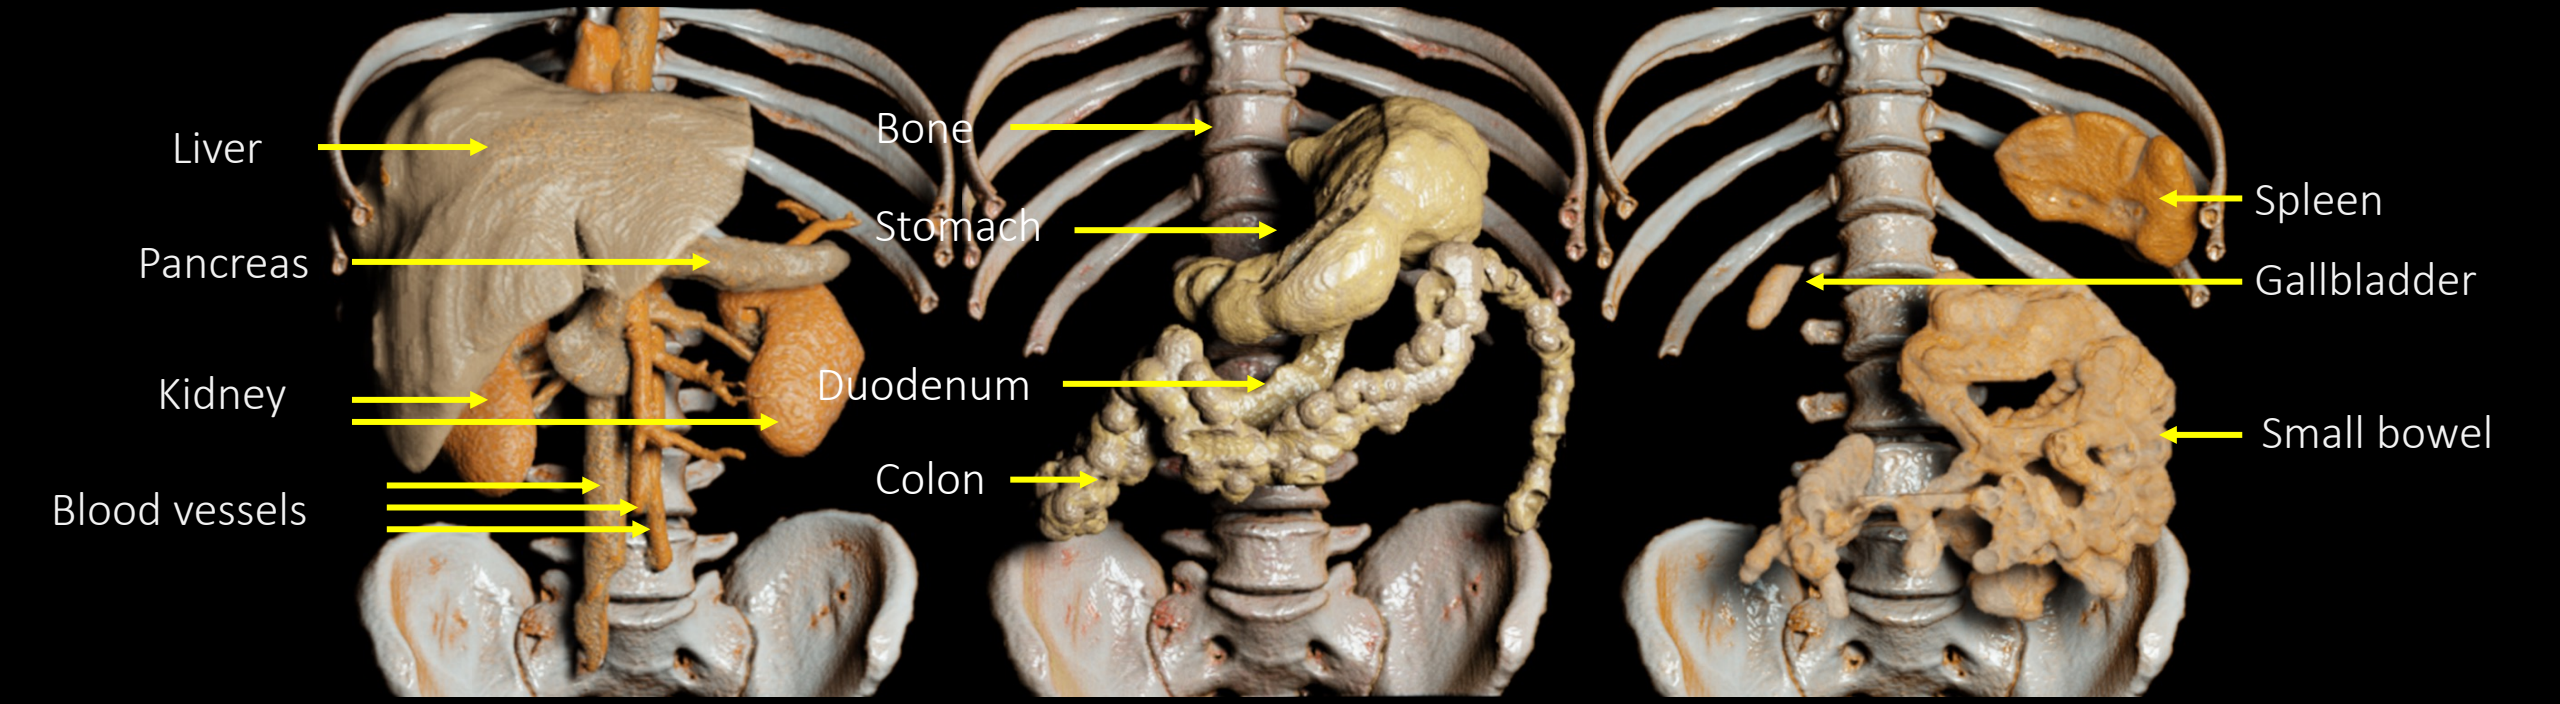

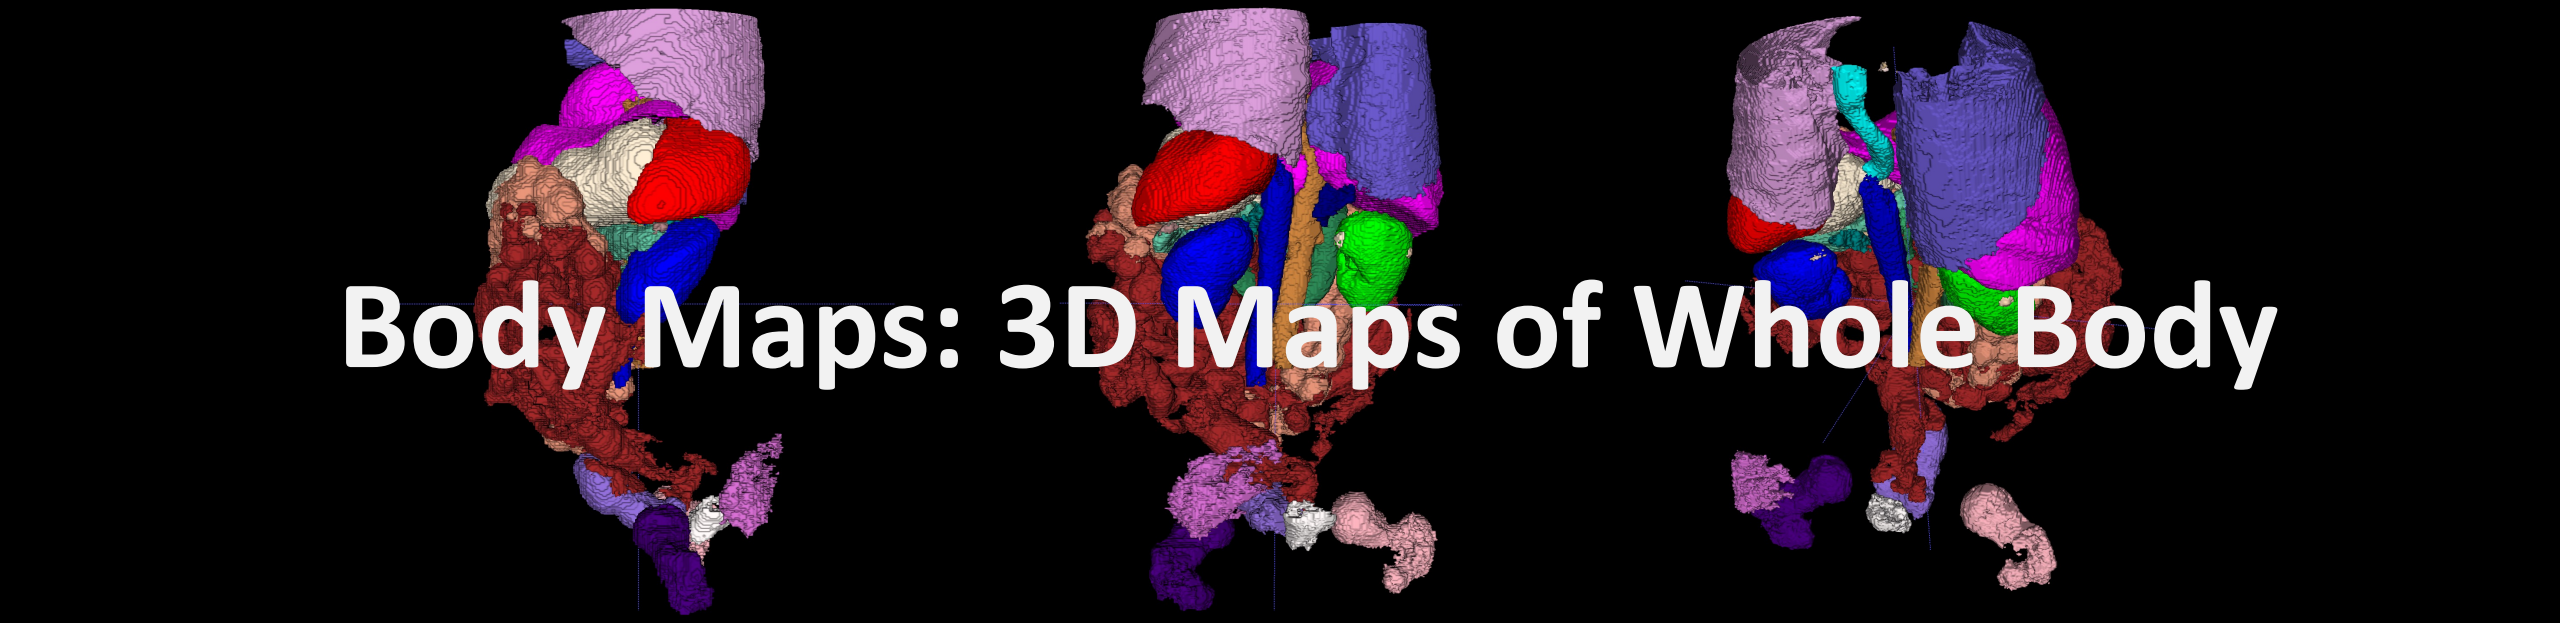

Body Maps: 3D Maps of Whole Body

Scaling Annotations, Datasets, and Algorithms for Medical Image Analysis

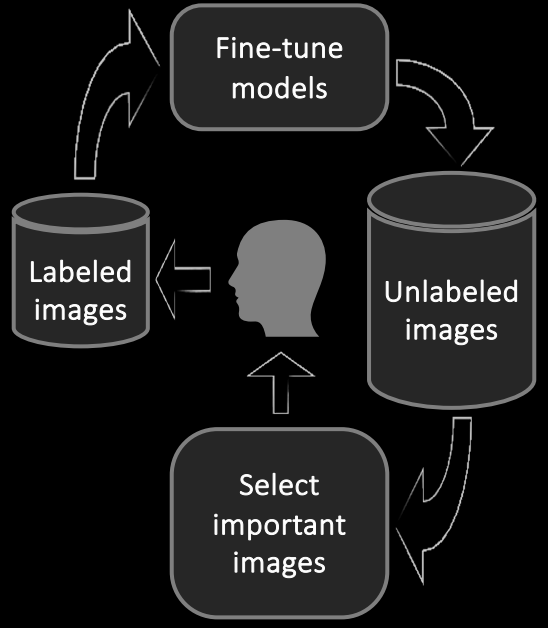

Interactive segmentation

Annotated
25
organs

Annotated
7
cancers

Integrated
15
public datasets

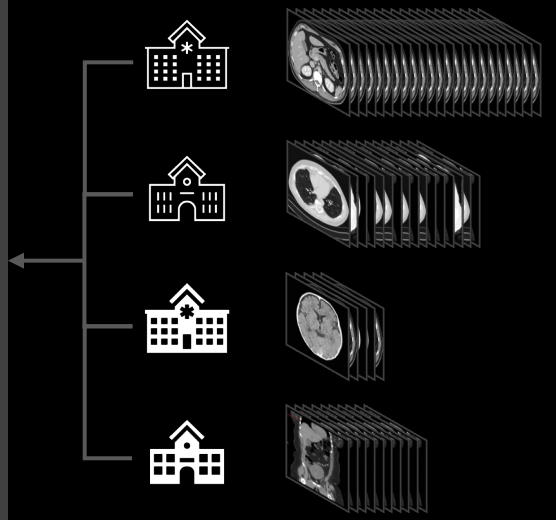

Collected from
27 hospitals
worldwide

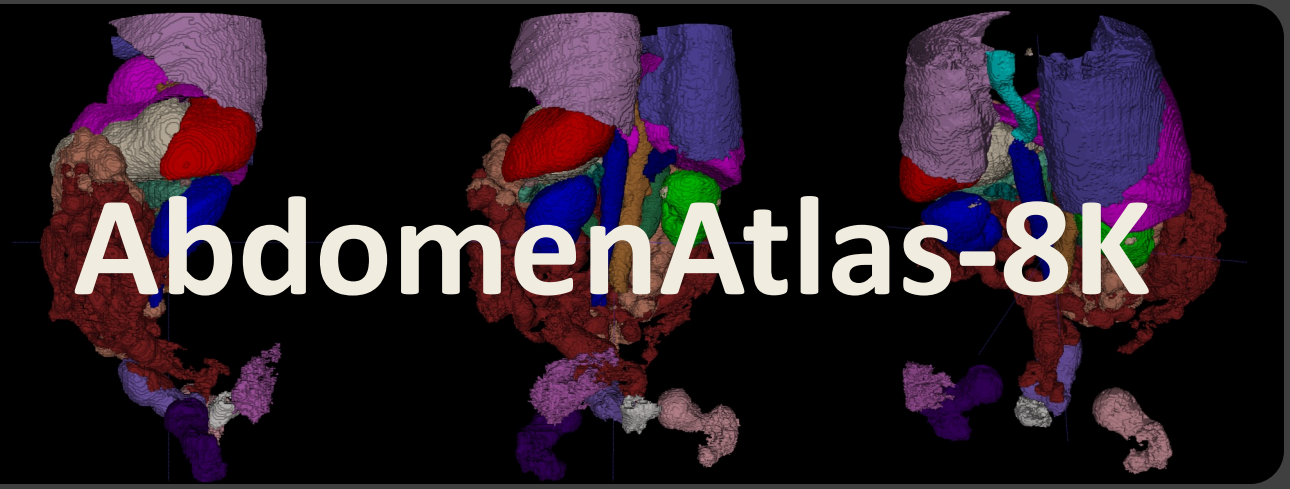

AbdomenAtlas-8K

Up to
533x faster
than previous strategies

MONAI

Annotated
3.2M
images

Annotated
8,448
CT volumes

Created in
3 Weeks
by 1 annotator

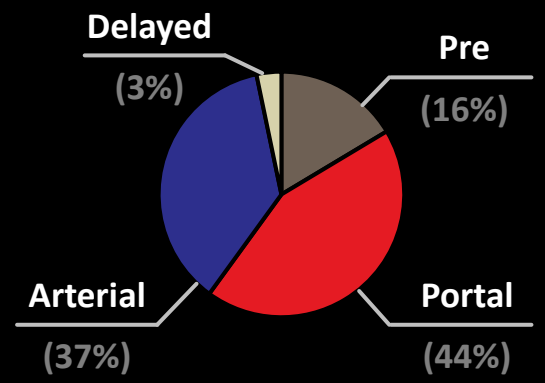

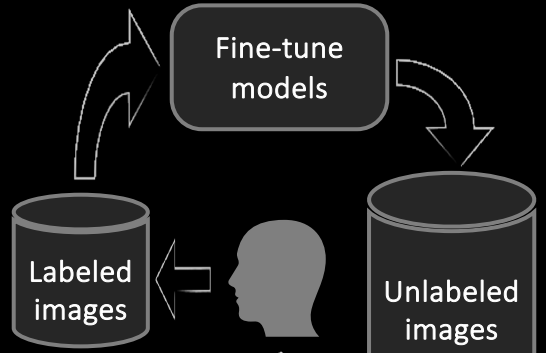

Interactive segmentation

Up to

533x faster

than previous strategies

MONAI

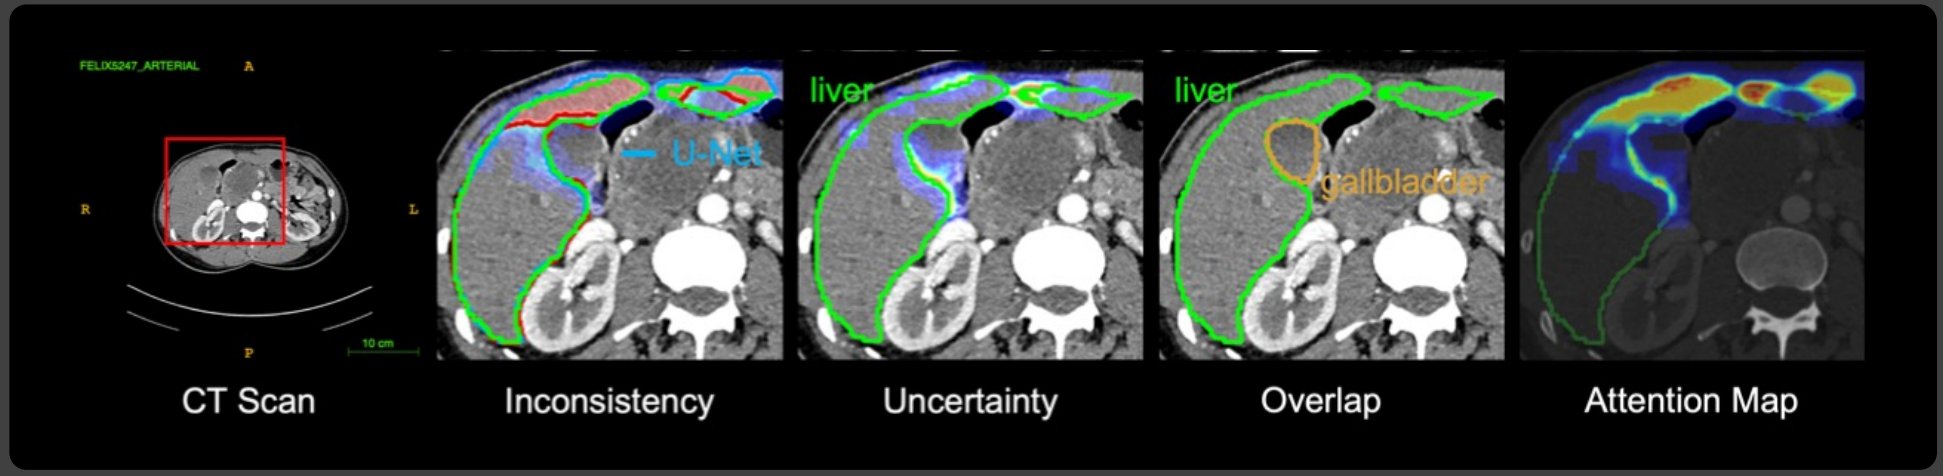

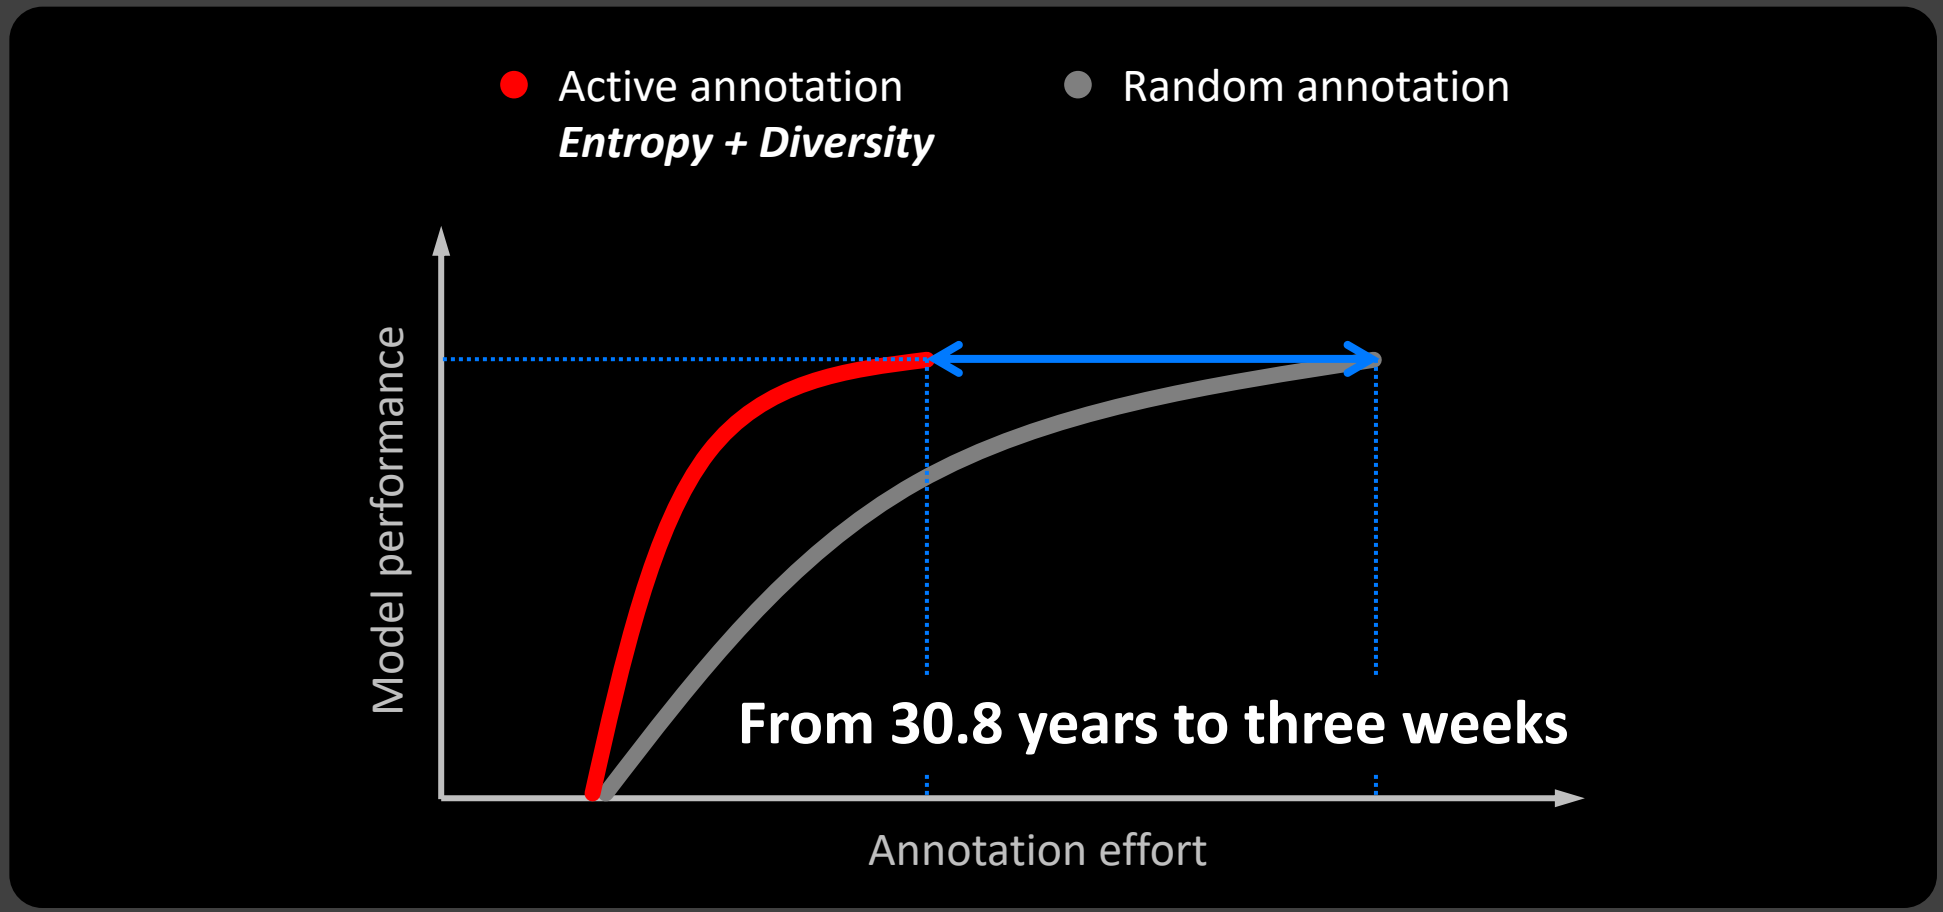

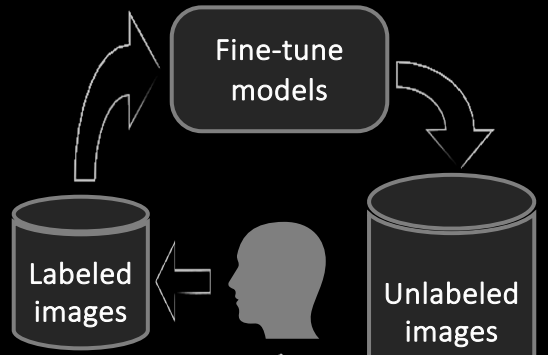

Interactive segmentation

Up to

533x faster

than previous strategies

MONAI

We will release **AbdomenAtlas-8K** of
8,448 CT volumes, totaling **3.2 million** CT slices

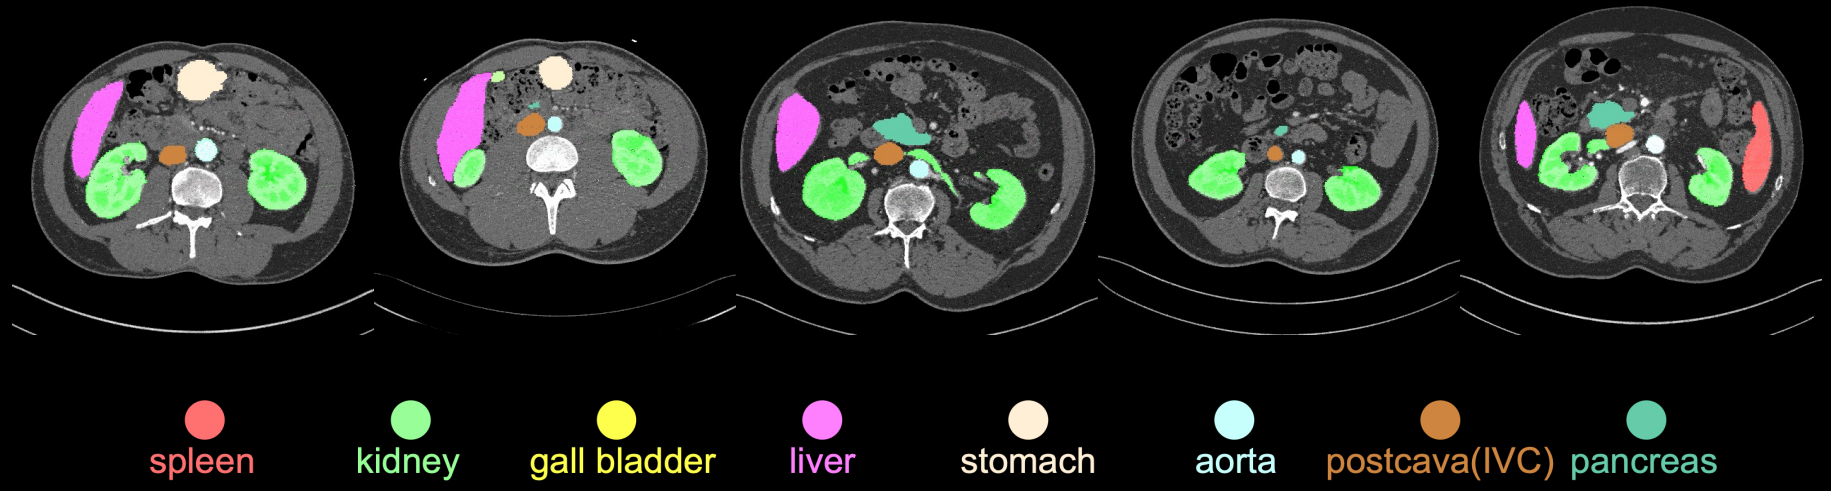

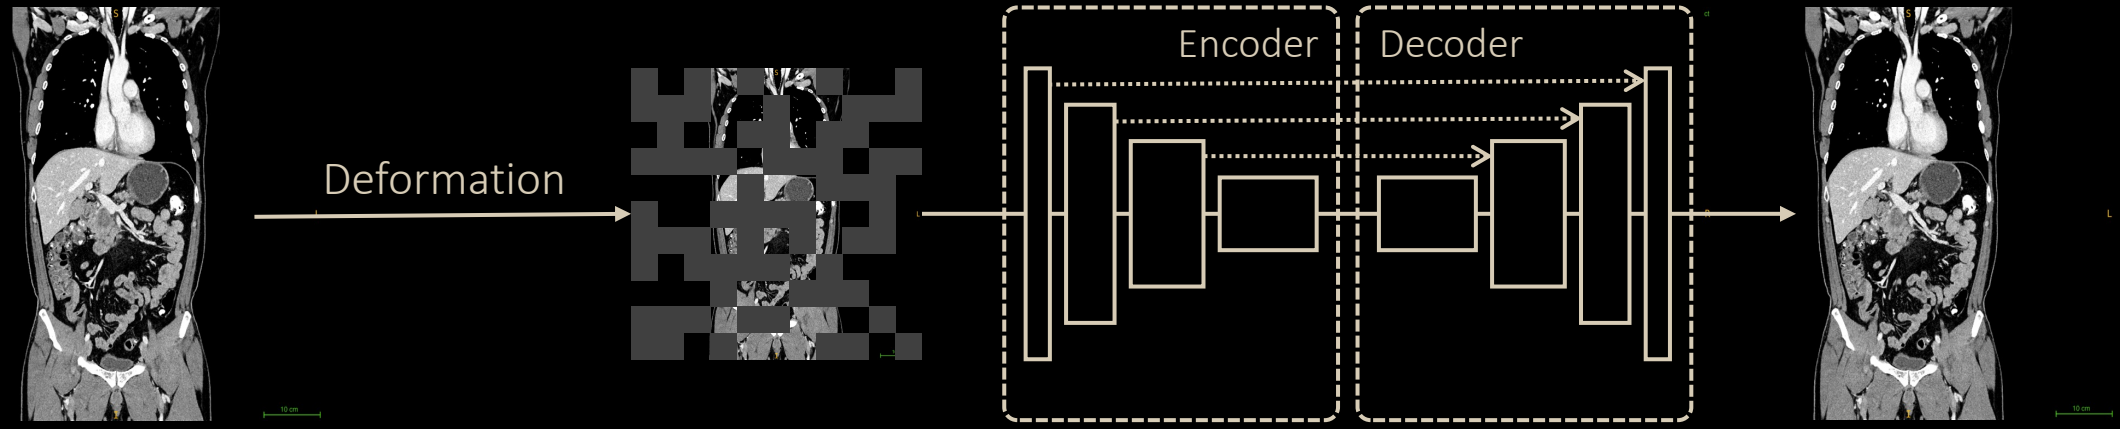

Self-supervised Pre-training

5,000 data + 1,152 GPU hours

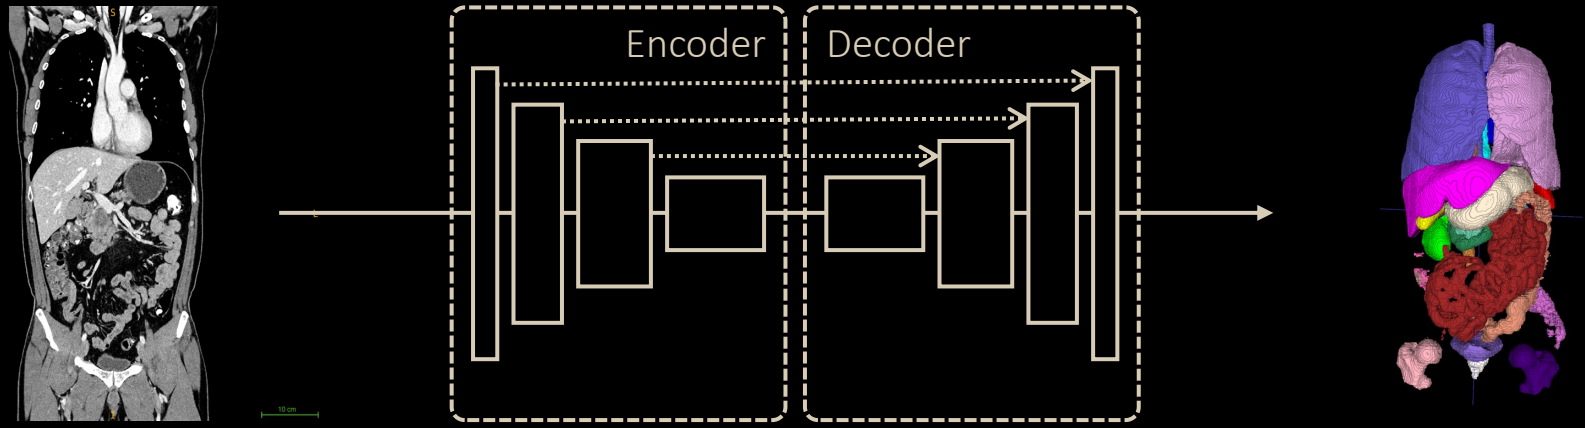

Supervised Pre-training

20 data + 20 annotation + 40 GPU hours

Scaling Annotations, Datasets, and Algorithms for Medical Image Analysis

Medical professionals cannot tell which are real and which are synthetic tumors

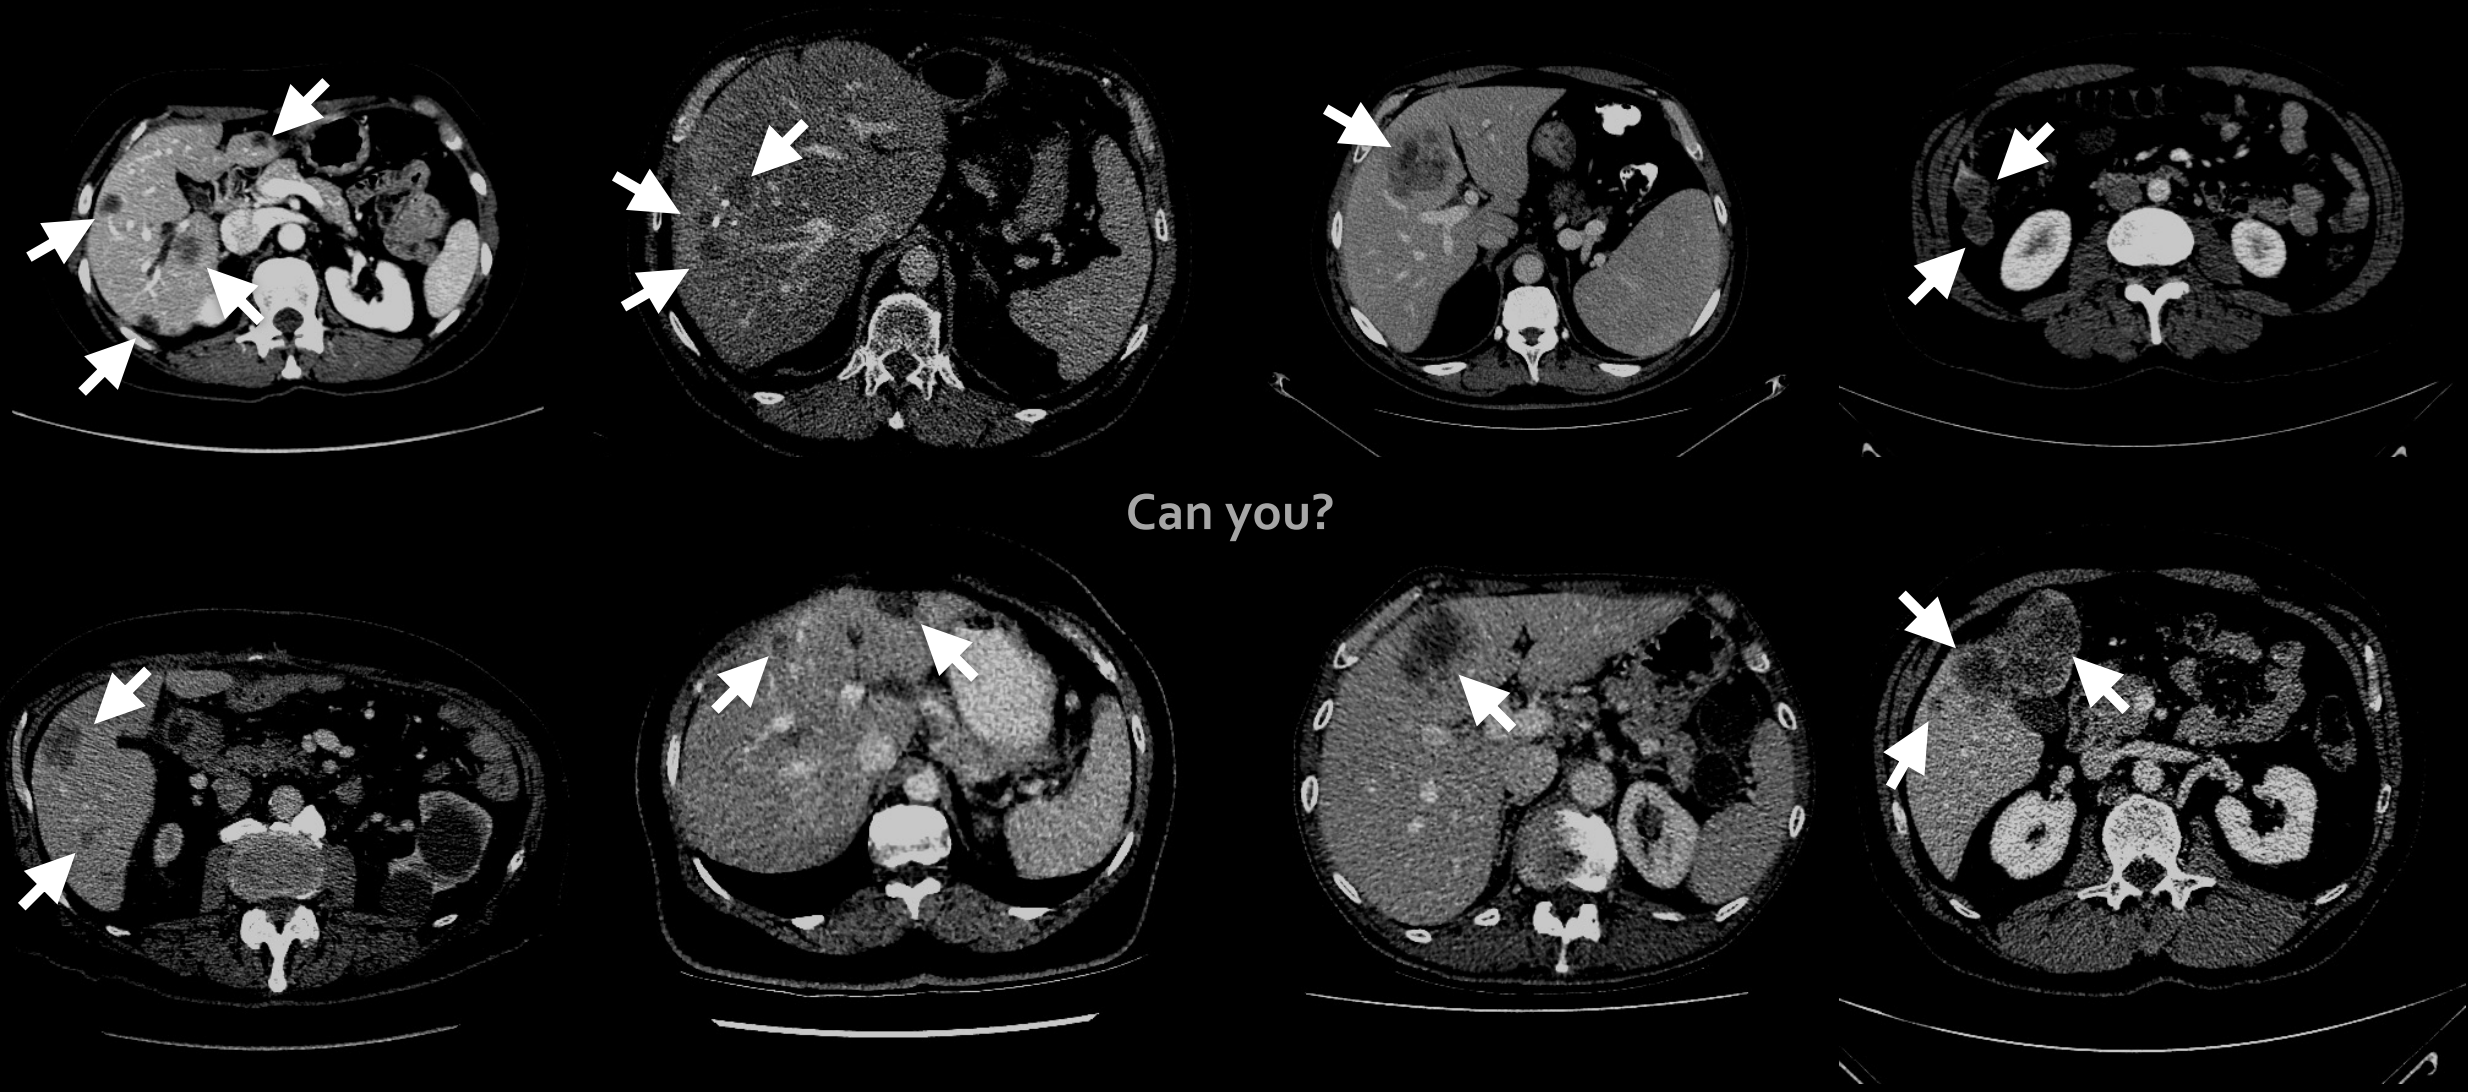

Can you?

Training AI on synthetic tumors performs as well as training it on real tumors

CT

AI prediction
trained on real tumors
with per-voxel annotation
DSC = 58% [52% - 63%]

AI prediction
trained on synthetic tumors
with no annotation
DSC = 60% [55% - 65%]

- Liver
- Liver tumor

[Qualitative] Generating enormous small tumors for training AI models

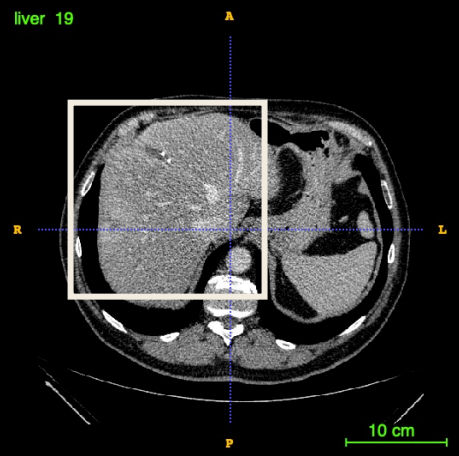

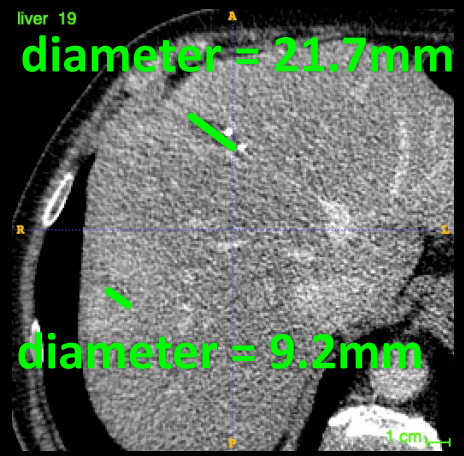

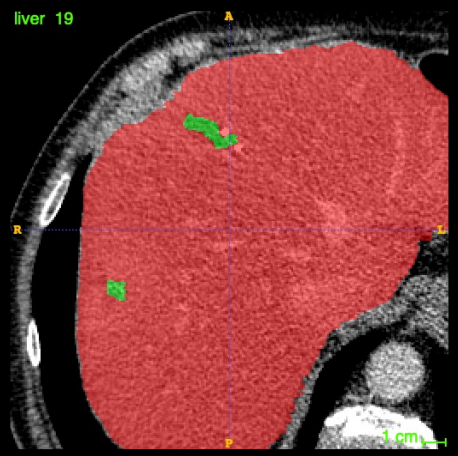

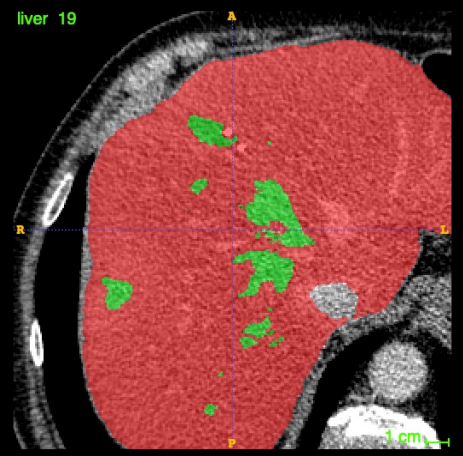

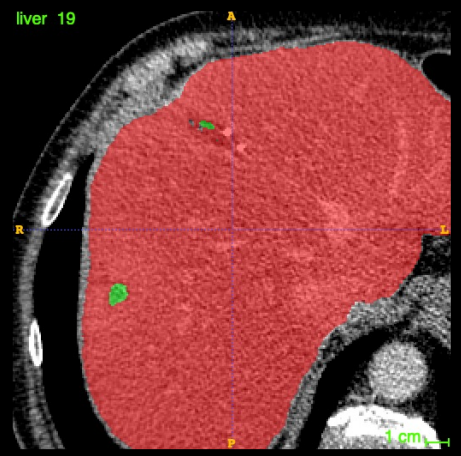

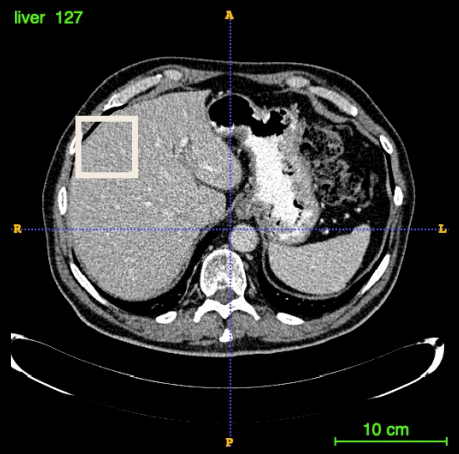

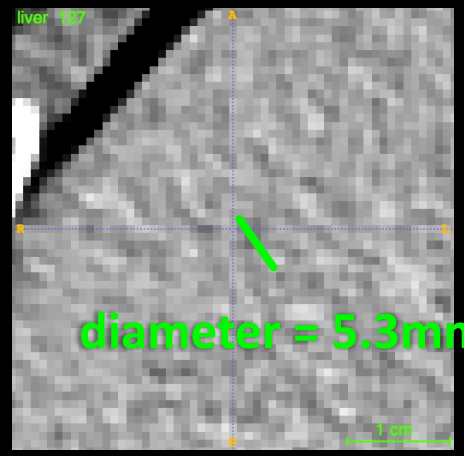

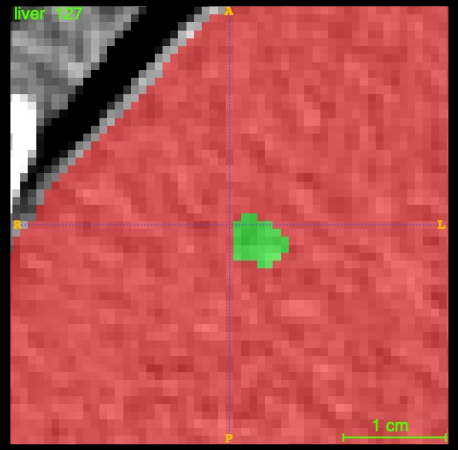

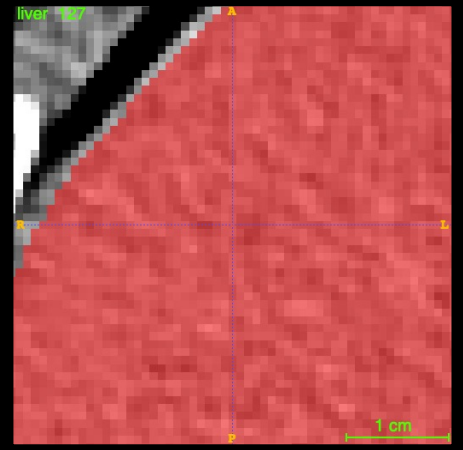

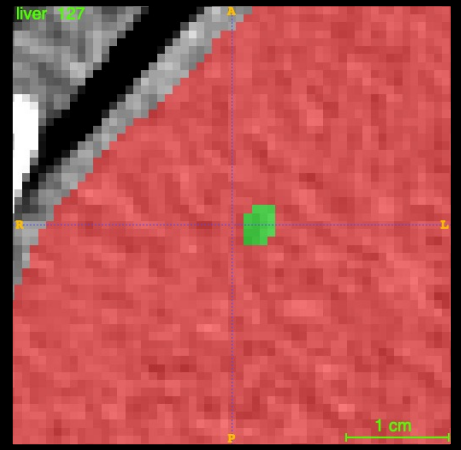

CT scan | zoom-in | ground truth | AI trained on real tumors | AI trained on synthetic tumors

[Quantitative] Generating enormous small tumors for training AI models

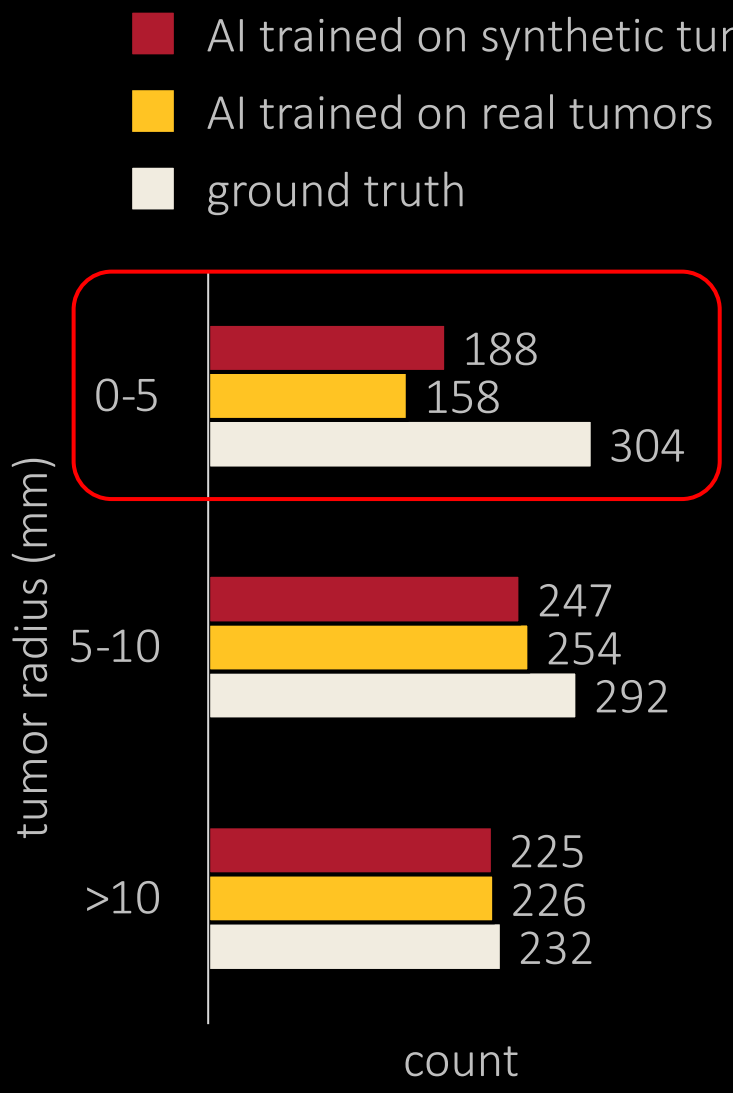

Observation: Compared with real tumors,
AI trained on synthetic tumors improves Sensitivity from 52% to 62% for detecting small tumors (0-5mm).

- Needed for early detection
  - Early signs of cancer can be subtle
  - 1/2 of liver cancer are missed by radiologists
- Needed for AI development
  - CT scans with early cancer are limited
  - Annotations for early cancer are hard
- Needed for medical education
  - Junior radiologists have an Accuracy of 20%
  - Senior radiologists have an Accuracy of 78%

Scaling Annotations, Datasets, and Algorithms for Medical Image Analysis

featured in

ChimeraX

at UCSF

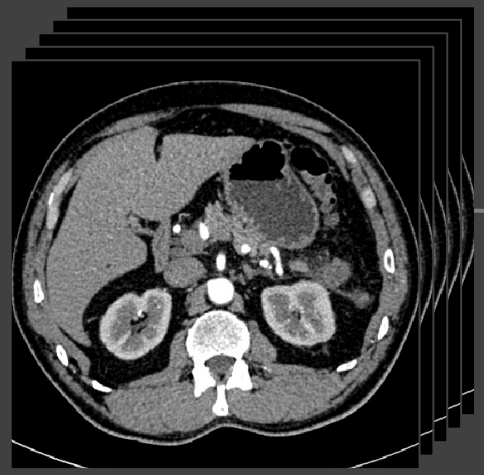

Vision Encoder

Universal Model

Please segment the tumor in the tail of the pancreas and then measure its size.

Take a look at these CT scans and mark the suspected tumor region.

.....

Text Encoder

featured in

MONAI

at NVIDIA

This tumor is likely to be PDAC with a diameter of 25mm.

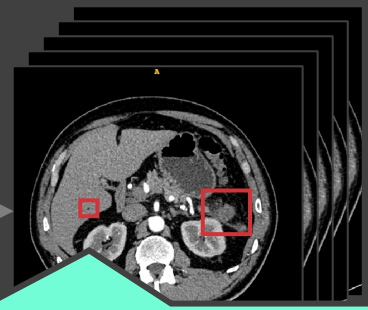

Two potential tumors are framed in bounding boxes.

Info | Teams | Submit | Leaderboard | Statistics

Challenge

Challenge Leaderboard

Search:

Additional metrics | Show all metrics

#	User (Team)	Created	Mean Position
1st	zongwei.zhou (universal_model)	13 Feb. 2023	5.6
2nd	Swin_UNETR	12 Nov. 2021	10.1
3rd	ahatamiz2	12 Nov. 2021	10.1
4th	Isensee	6 Dec. 2019	12.3
5th	AndyL	24 Nov. 2022	12.5
6th	heyufan1995	30 Oct. 2020	12.8
7th	qsyeung	5 Jan. 2023	12.9
8th	vishwesh.nath	11 Nov. 2021	13.5

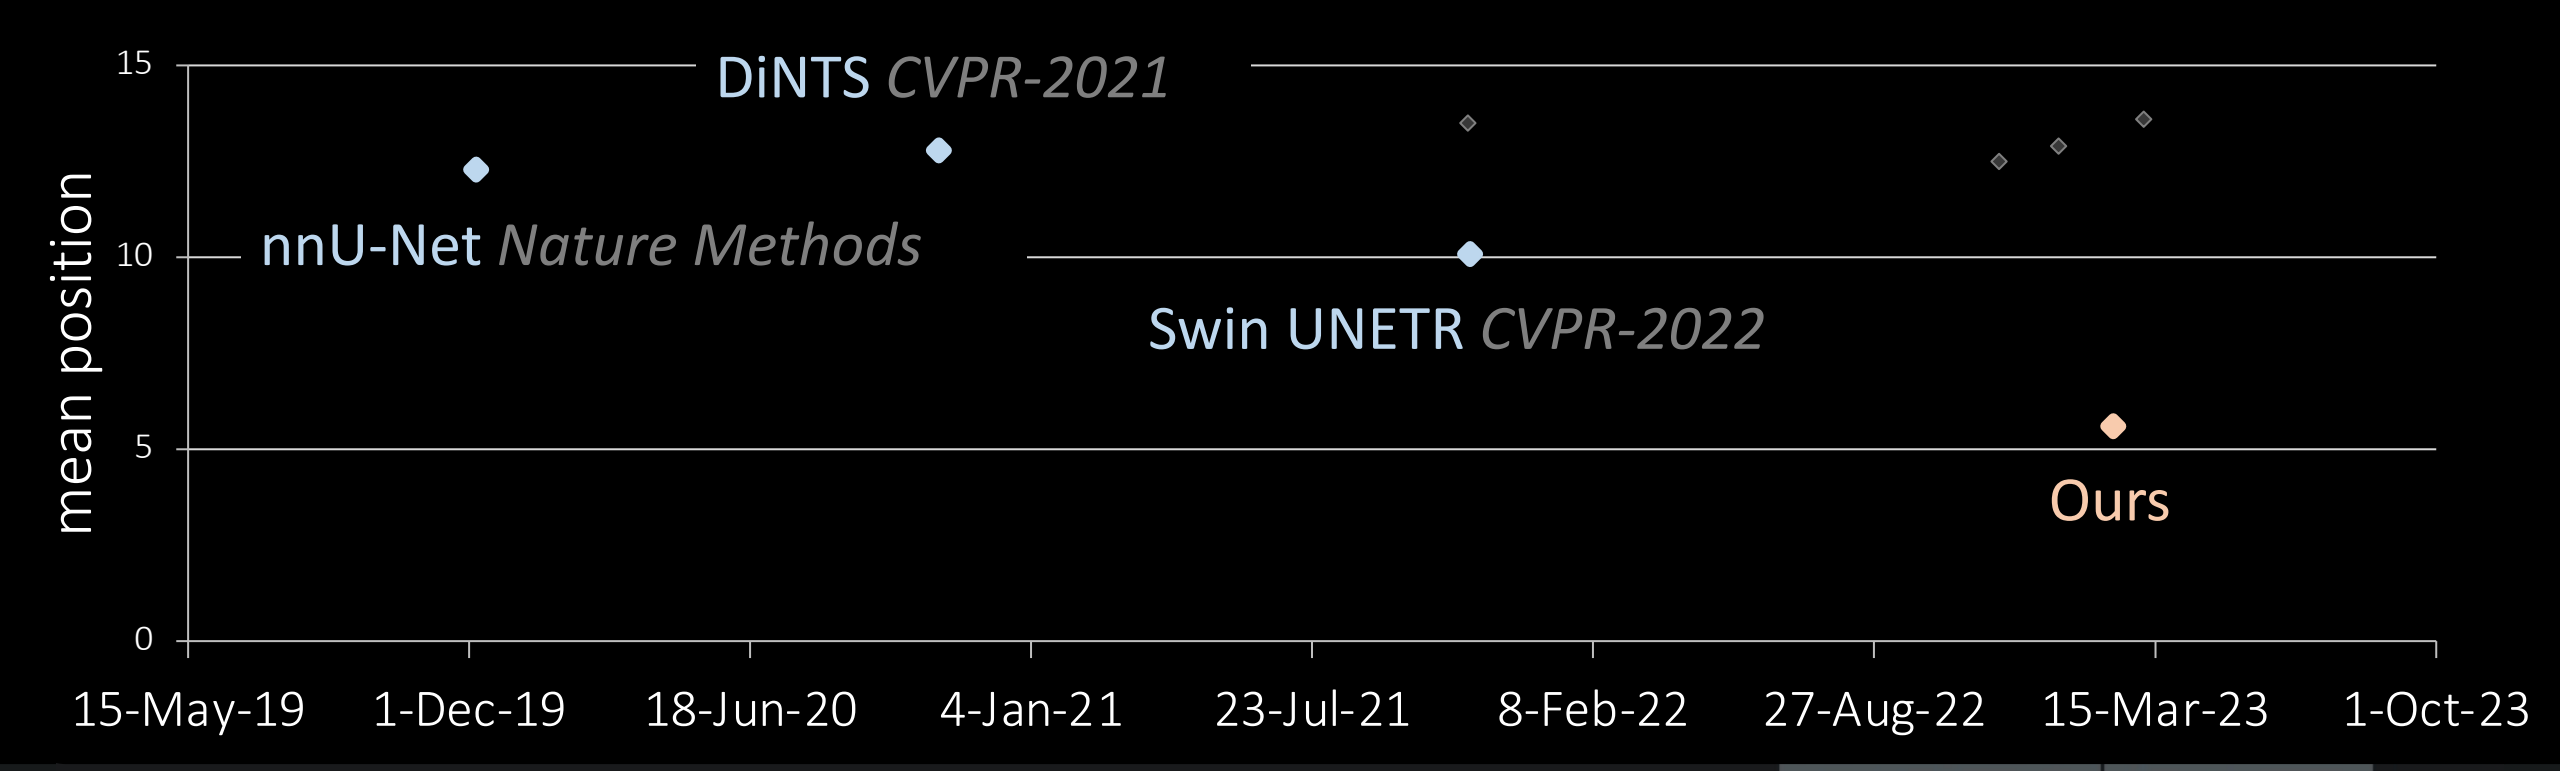

Additional metrics | Show all metrics

#	User (Team)	Created	Mean Position
1st	zongwei.zhou (universal_model)	13 Feb. 2023	5.6
2nd	Swin_UNETR	12 Nov. 2021	10.1
3rd	ahatamiz2	12 Nov. 2021	10.1
4th	Isensee	6 Dec. 2019	12.3
5th	AndyL	24 Nov. 2022	12.5
6th	heyufan1995	30 Oct. 2020	12.8
7th	qsyeung	5 Jan. 2023	12.9
8th	vishwesh.nath	11 Nov. 2021	13.5

The AI predictions for six organs are comparable to expert annotators

If we spend a lot more money to ask radiologists to annotate these six organs, it might turn out that the AI can do a similar quality annotation

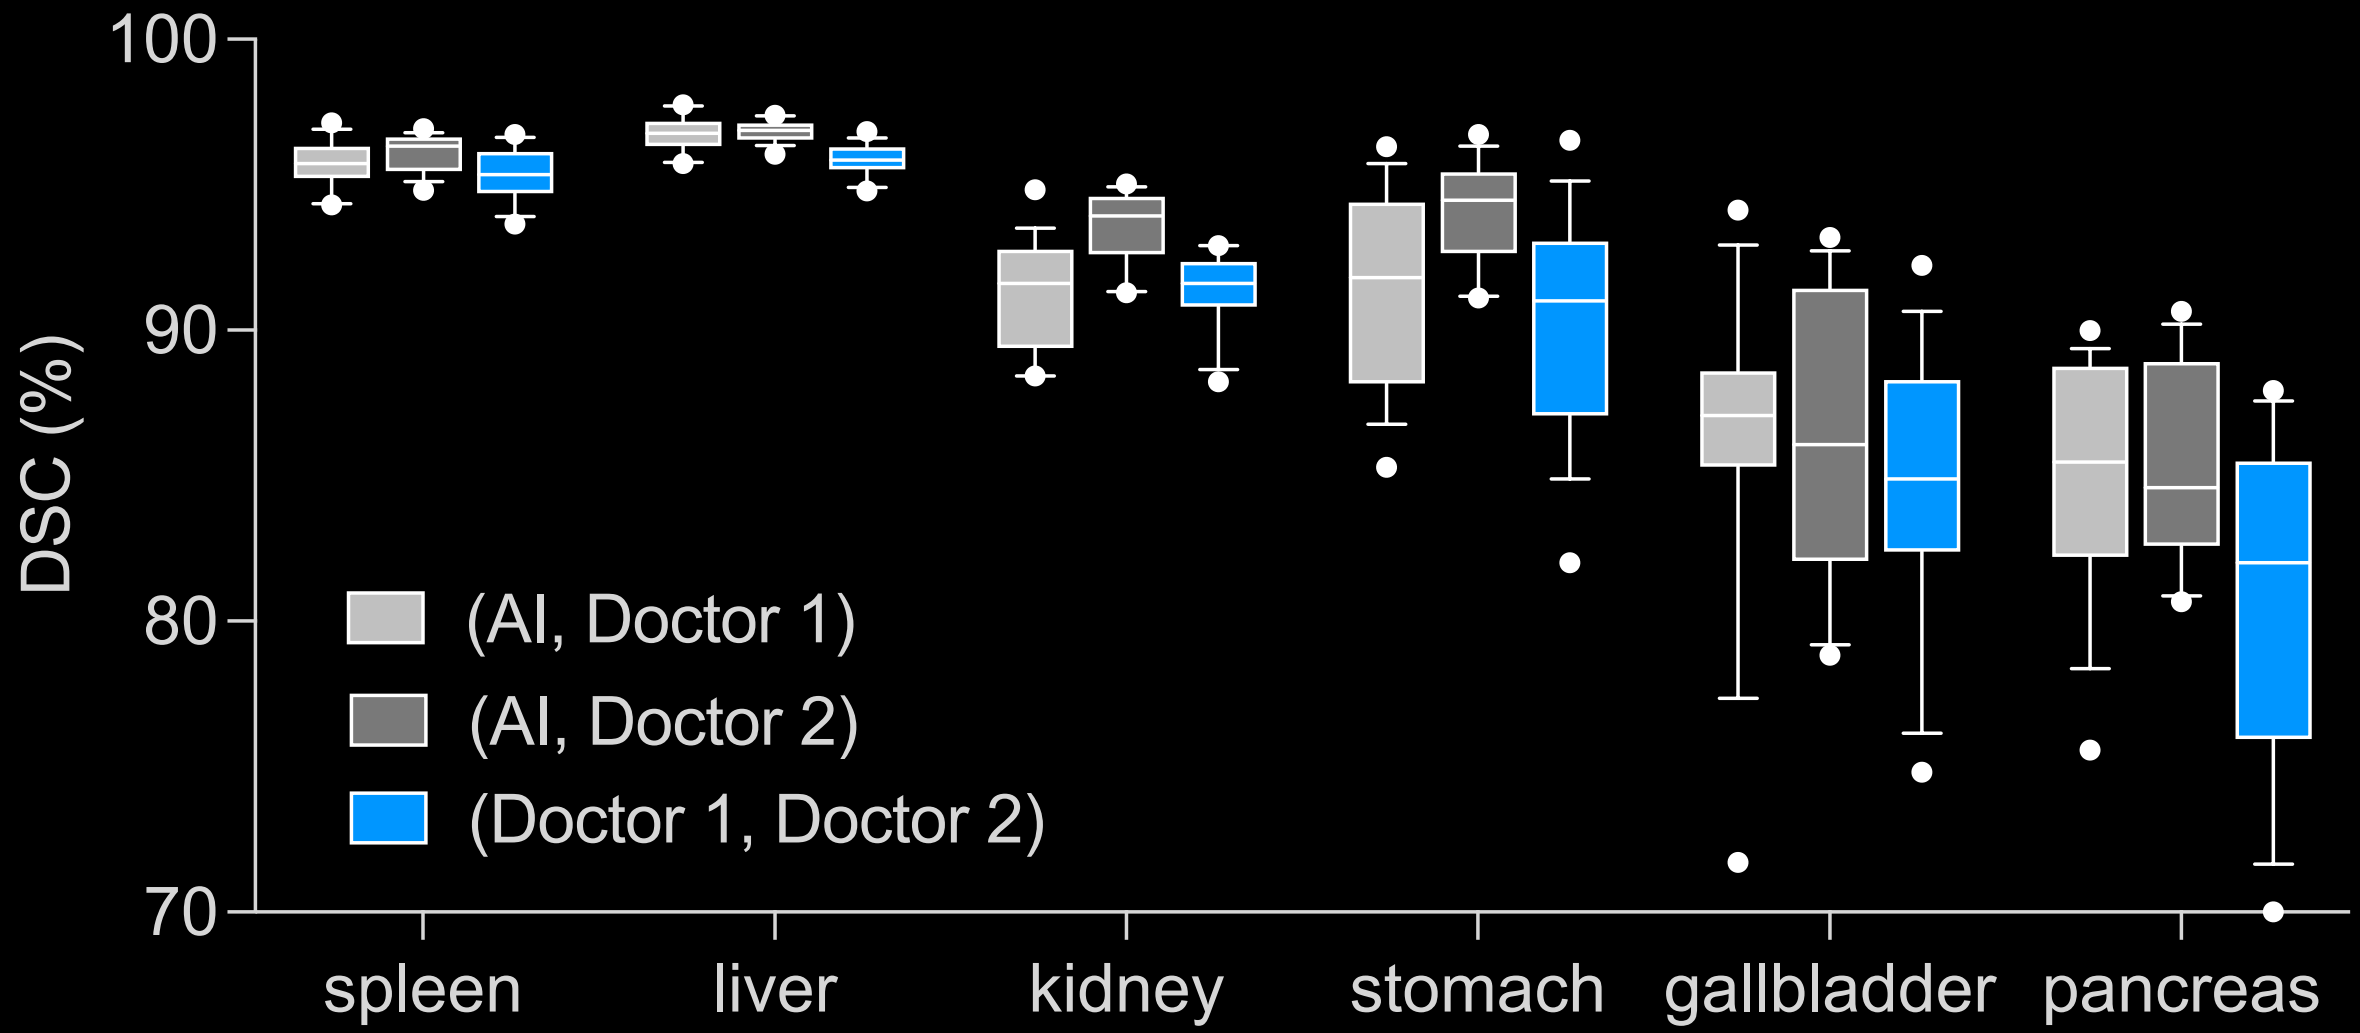

The AI predictions for six organs are comparable to expert annotators

If we spend a lot more money to ask radiologists to annotate these six organs, it might turn out that the AI can do a similar quality annotation

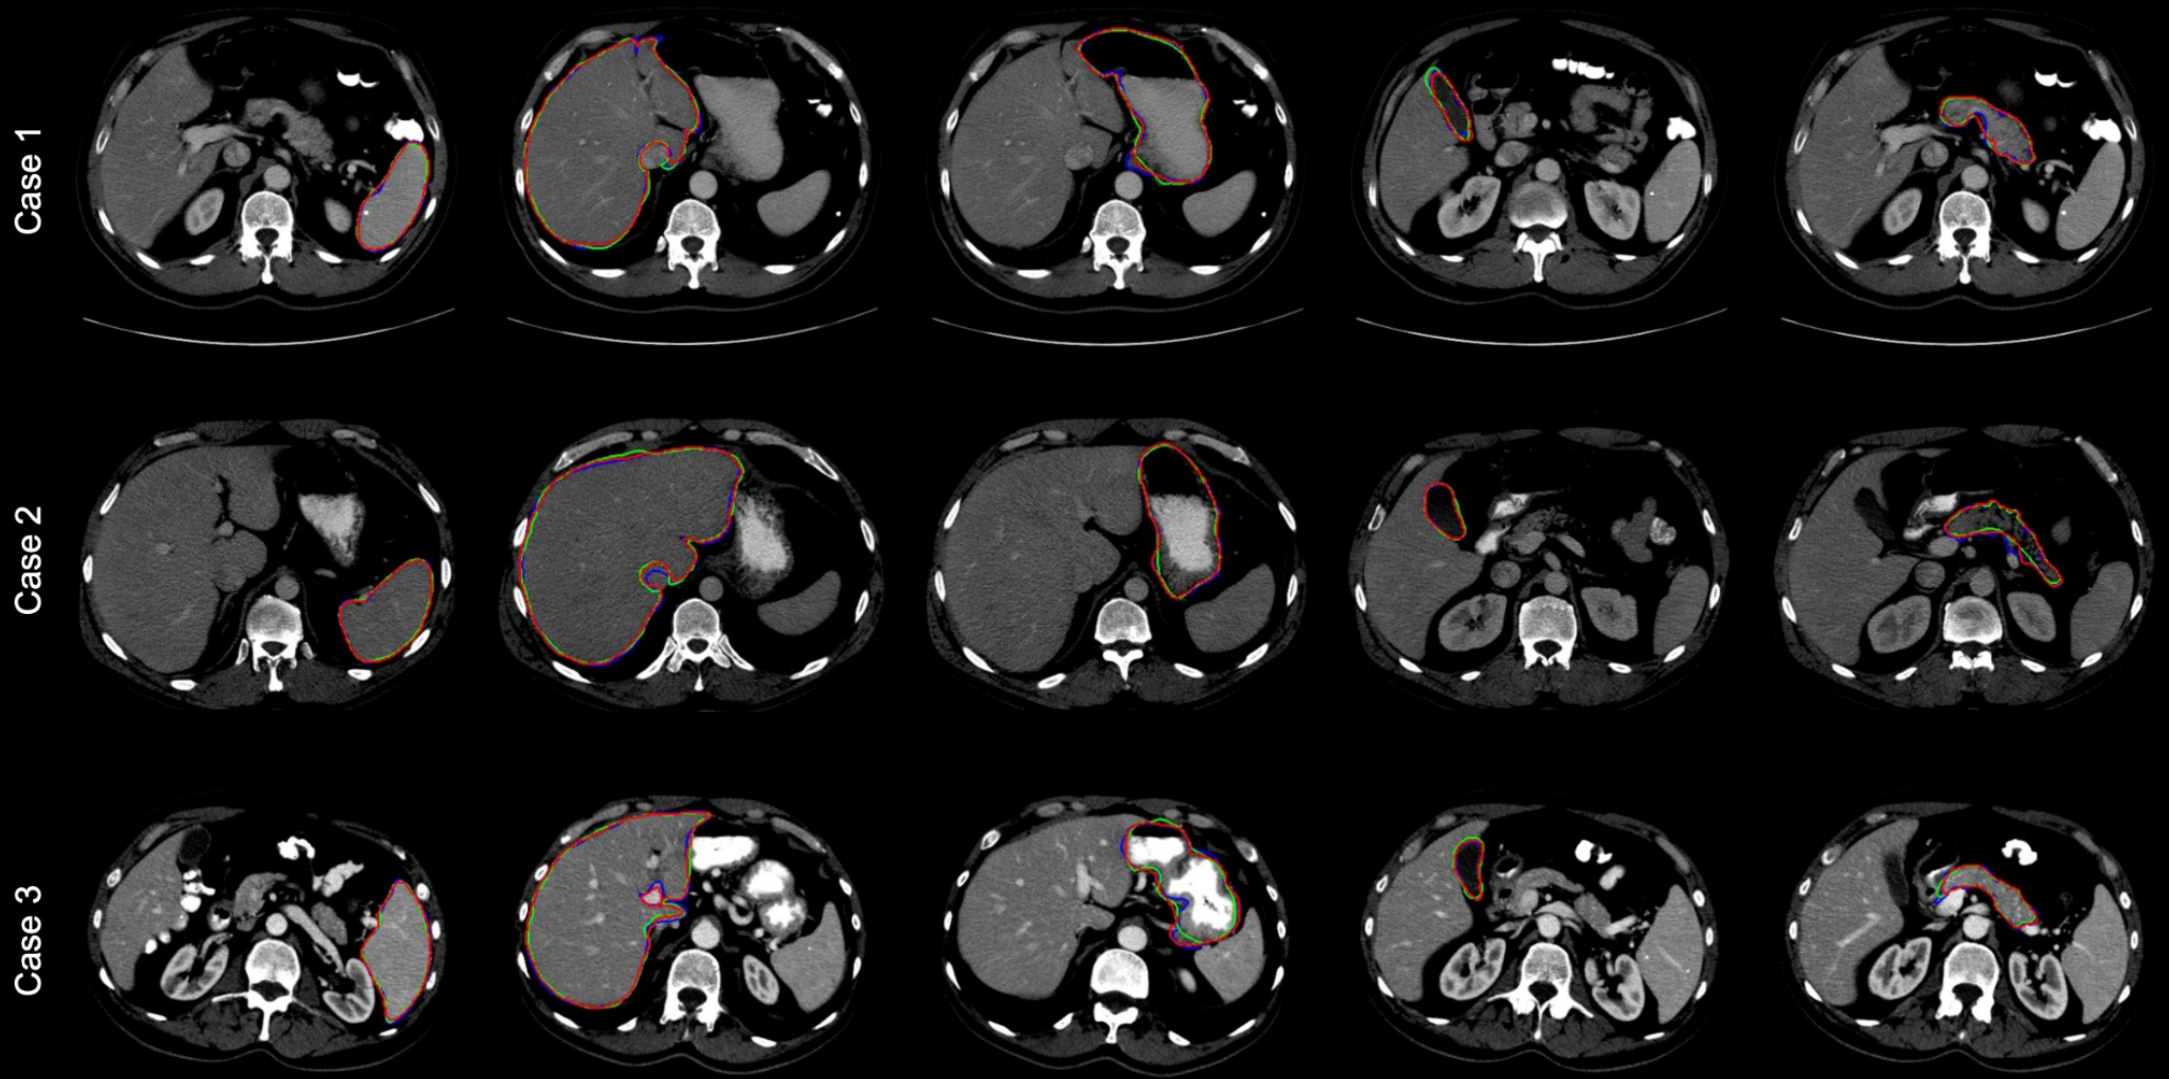

Scaling annotations

Efficient annotation
Human in the loop
Novel disease

AbdomenAtlas-8K

8,448 annotated CT volumes

Scaling datasets

Multiple modalities
Diverse institutes
IRB approval

Tumor Synthesis

Annotation-free deep learning

Scaling algorithms

Vision-language
Lifelong learning
Reader study

Universal Model

25 organs and 7 cancers

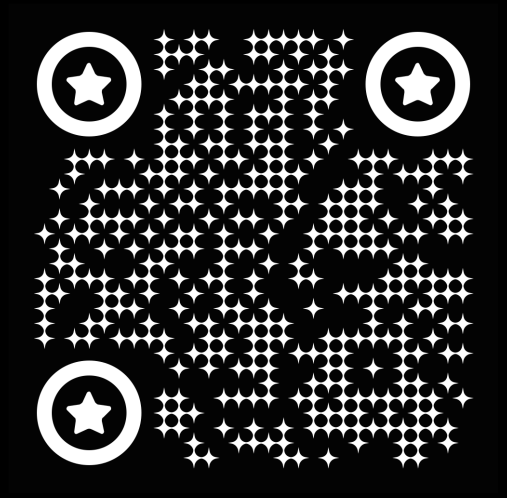

Code & Dataset
NeurIPS-2023

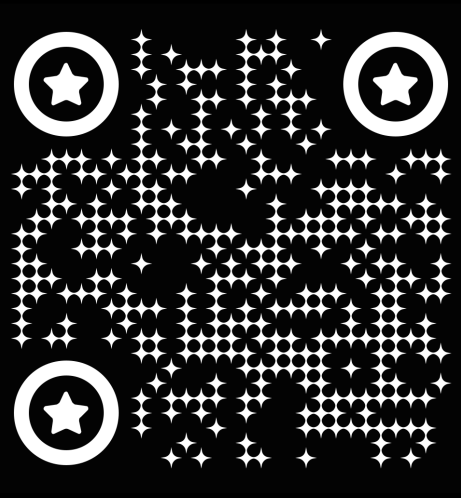

Code & Turing Test
CVPR-2023

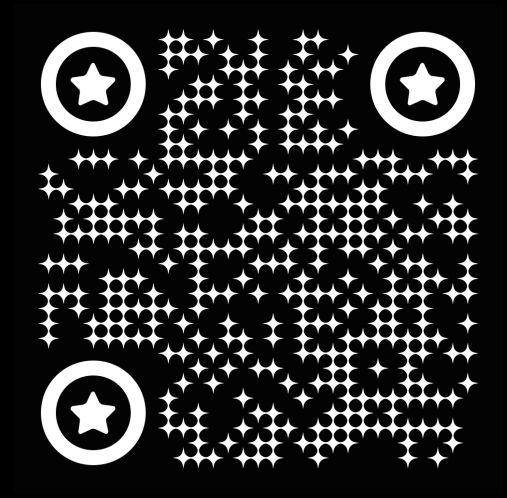

Code & Model
ICCV-2023

Reference

Scaling Annotations

C. Qu, T. Zhang, H. Qiao, J. Liu, Y. Tang, A. Yuille, and Z. Zhou*. "Annotating 8,000 Abdominal CT Volumes for Multi-Organ Segmentation in Three Weeks." NeurIPS 2023.

Z. Zhou, J. Shin, L. Zhang, S. Gurudu, M. Gotway, and J. Liang*. “Fine-tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally.” CVPR 2017.

Z. Zhou, J. Shin, S. Gurudu, M. Gotway, and J. Liang*. “Active, Continual Fine Tuning of Convolutional Neural Networks for Reducing Annotation Efforts.” Medical Image Analysis 2021.

Scaling Datasets

Q. Hu, Y. Chen, J. Xiao, S. Sun, J. Chen, A. Yuille, Z. Zhou*. “Label-Free Liver Tumor Segmentation.” CVPR 2023.

B. Li, Y. Chou, S. Sun, H. Qiao, A. Yuille, Z. Zhou*. “Early Detection and Localization of Pancreatic Cancer by Label-Free Tumor Synthesis.” MICCAI Workshop 2023.

Scaling Algorithms

J. Liu, Y. Zhang, J. Chen, Y. Lu, Y. Yuan, A. Yuille, Y. Tang*, Z. Zhou*. “CLIP-Driven Universal Model for Organ Segmentation and Tumor Detection.” ICCV 2023.

Y. Zhang, X. Li, H. Chen, A. Yuille, Y. Liu*, Z. Zhou*. “Learning without Forgetting for Continual Abdominal Multi-Organ and Tumor Segmentation.” MICCAI 2023.

M. Kang, B. Li, Z. Zhu, Y. Lu, E. Fishman, A. Yuille, Z. Zhou*. “Label-Assemble: Leveraging Multiple Datasets with Partial Labels.” ISBI 2023.