Scaling Datasets, Annotations, 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
www.zongweiz.com

Statistics

7,118 citations

Top 2% of Scientists in 2022

24 first/corresponding authored papers

Significant

Needful

Applications

Methodologies

Innovative

Impactful

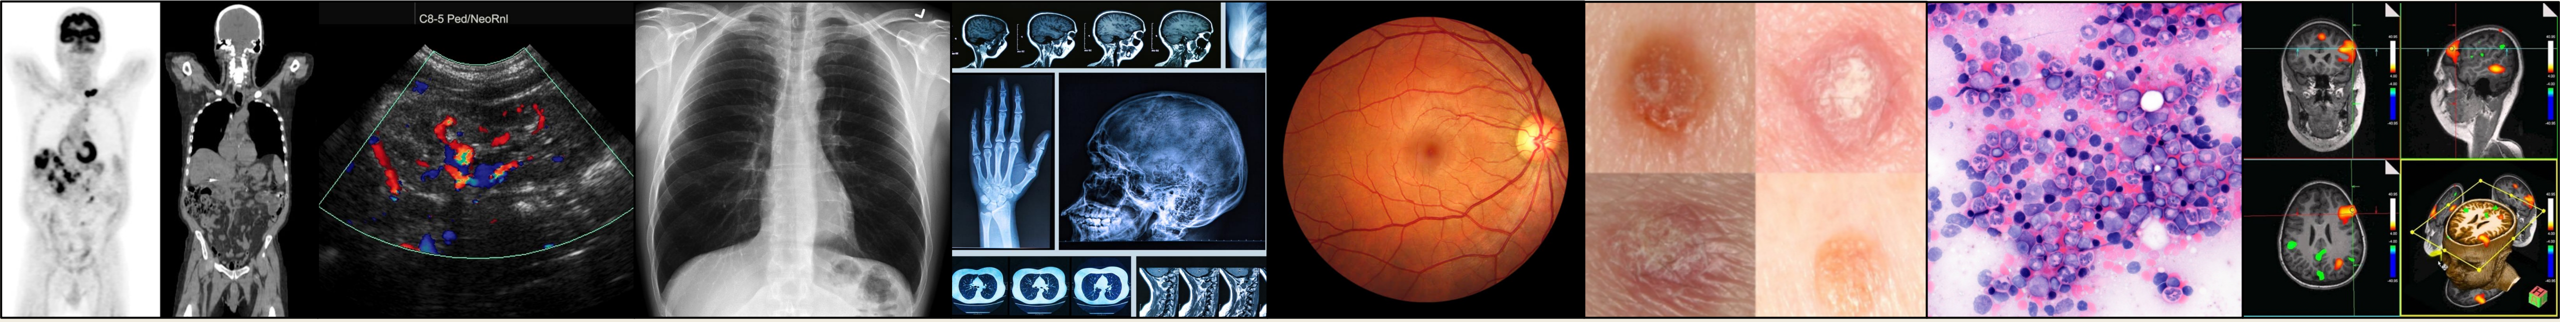

Imaging data account for about **90%** of all healthcare data

Significant

Applications

Methodologies

Impactful

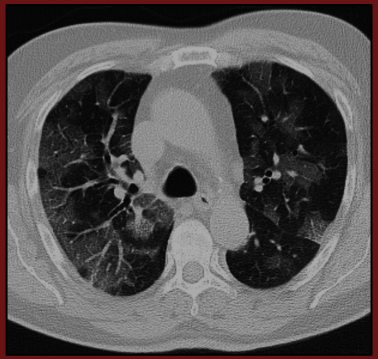

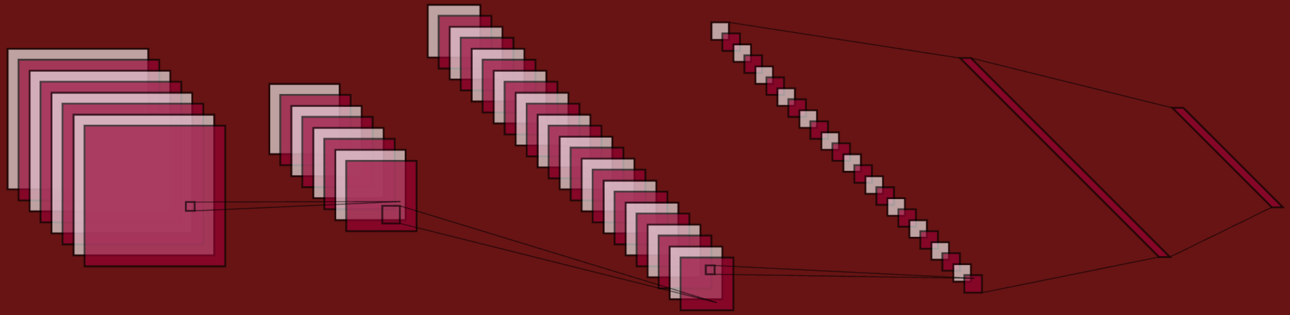

Lung cancer?

Malignant

Benign

Input image

Hidden layers

Output

Deep Learning has ushered in a revolution in medical imaging

CT

Ground Truth

annotated by human experts

Deep Learning

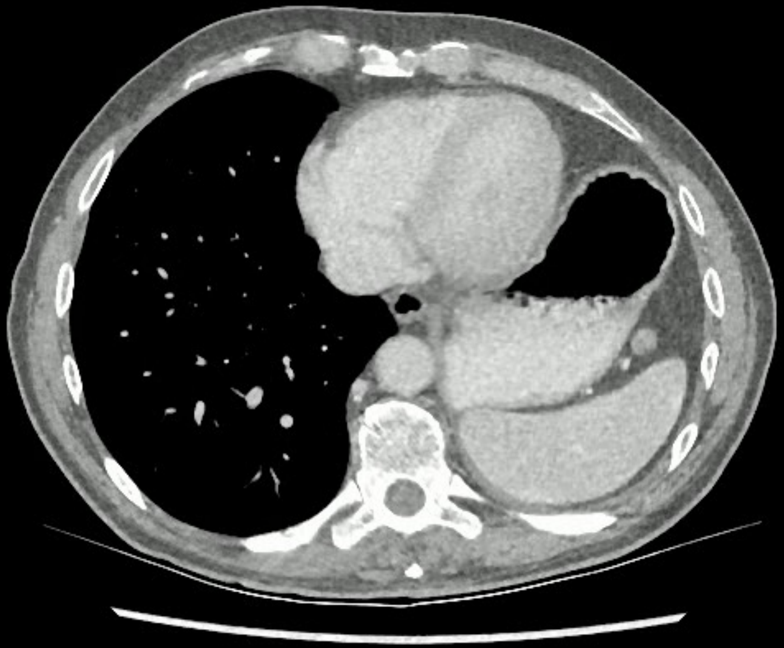

Liver

Liver tumor

1. Zhou, Zongwei. "Towards annotation-efficient deep learning for computer-aided diagnosis." PhD diss., Arizona State University, 2021.

Radiologists hate annotation, but computer scientists love annotation.

CT

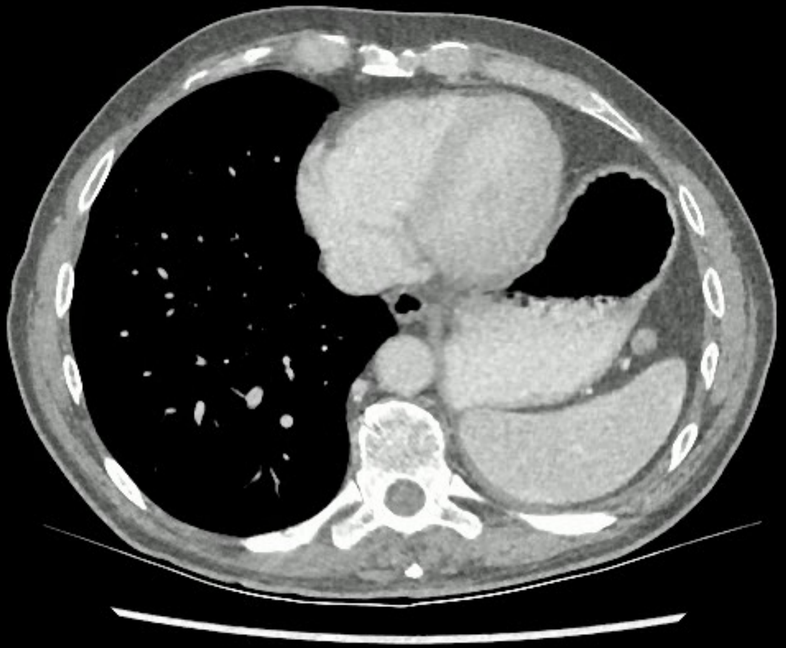

Ground Truth

annotated by human experts

Deep Learning

Liver

Liver tumor

Not enough annotation

1. Zhou, Zongwei. "Towards annotation-efficient deep learning for computer-aided diagnosis." PhD diss., Arizona State University, 2021.

Chapter I (2016-2022)

Methodologies: Reducing Annotation Efforts for Radiologists

Significant

Applications

Methodologies

Impactful

Computer-*Aided* Diagnosis

Assisting expert radiologists to see more patients and
to deliver more accurate diagnosis (*beyond human eye*)

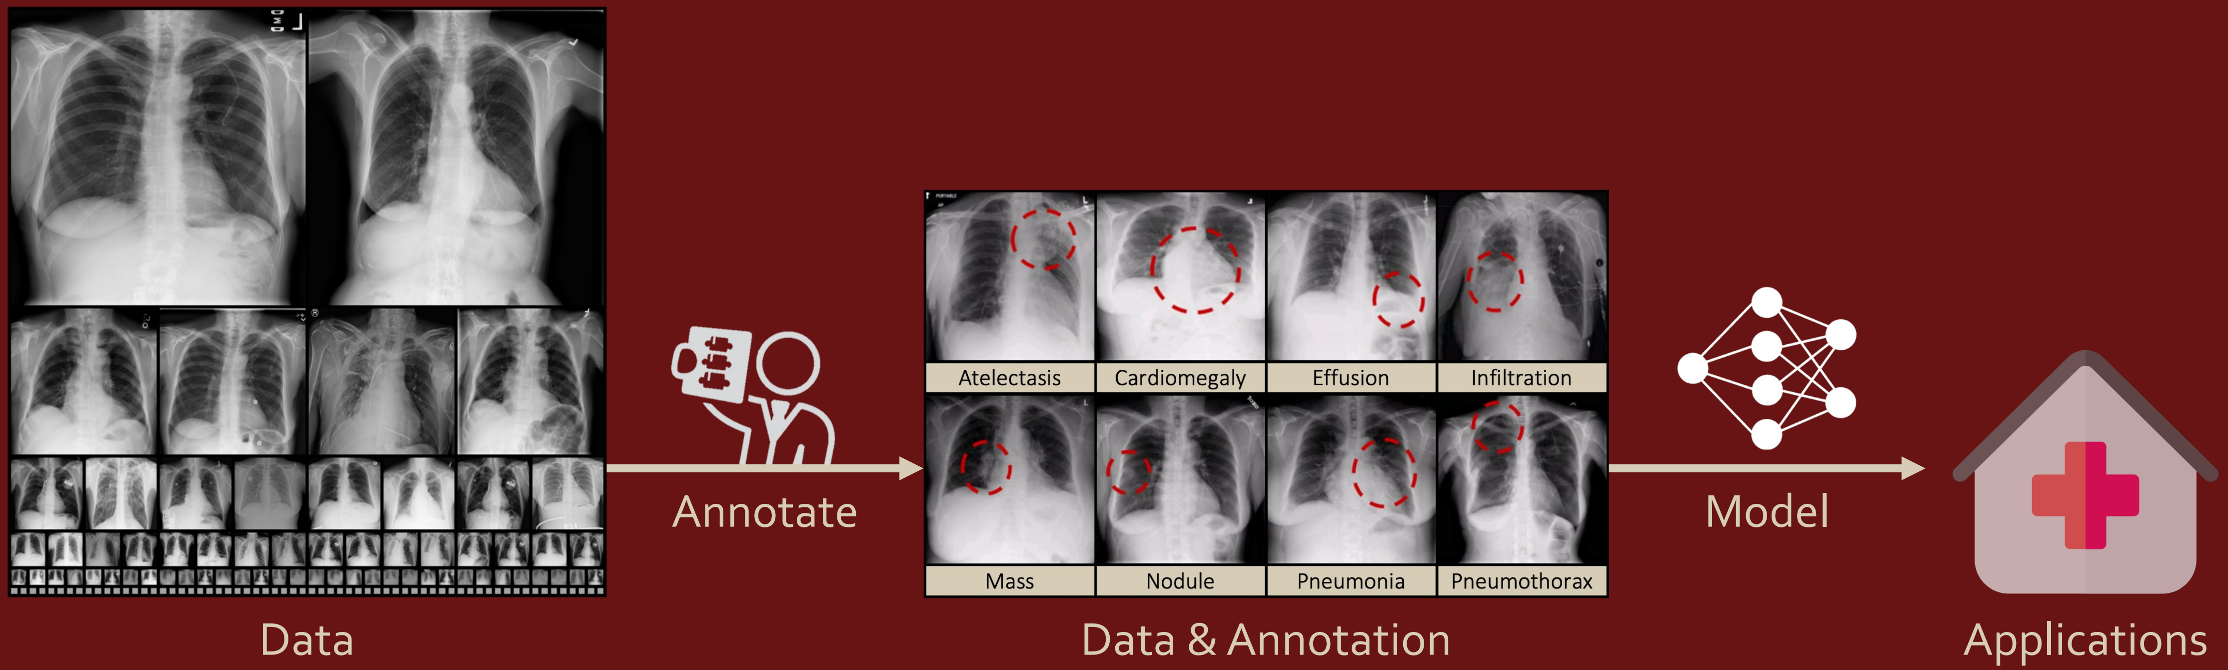

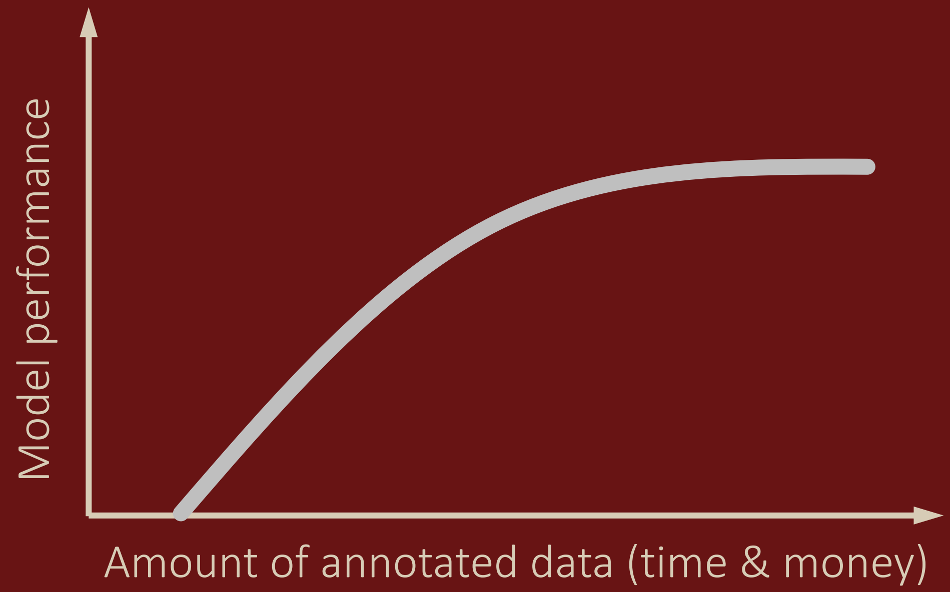
Model performance
Amount of annotated data (time & money)

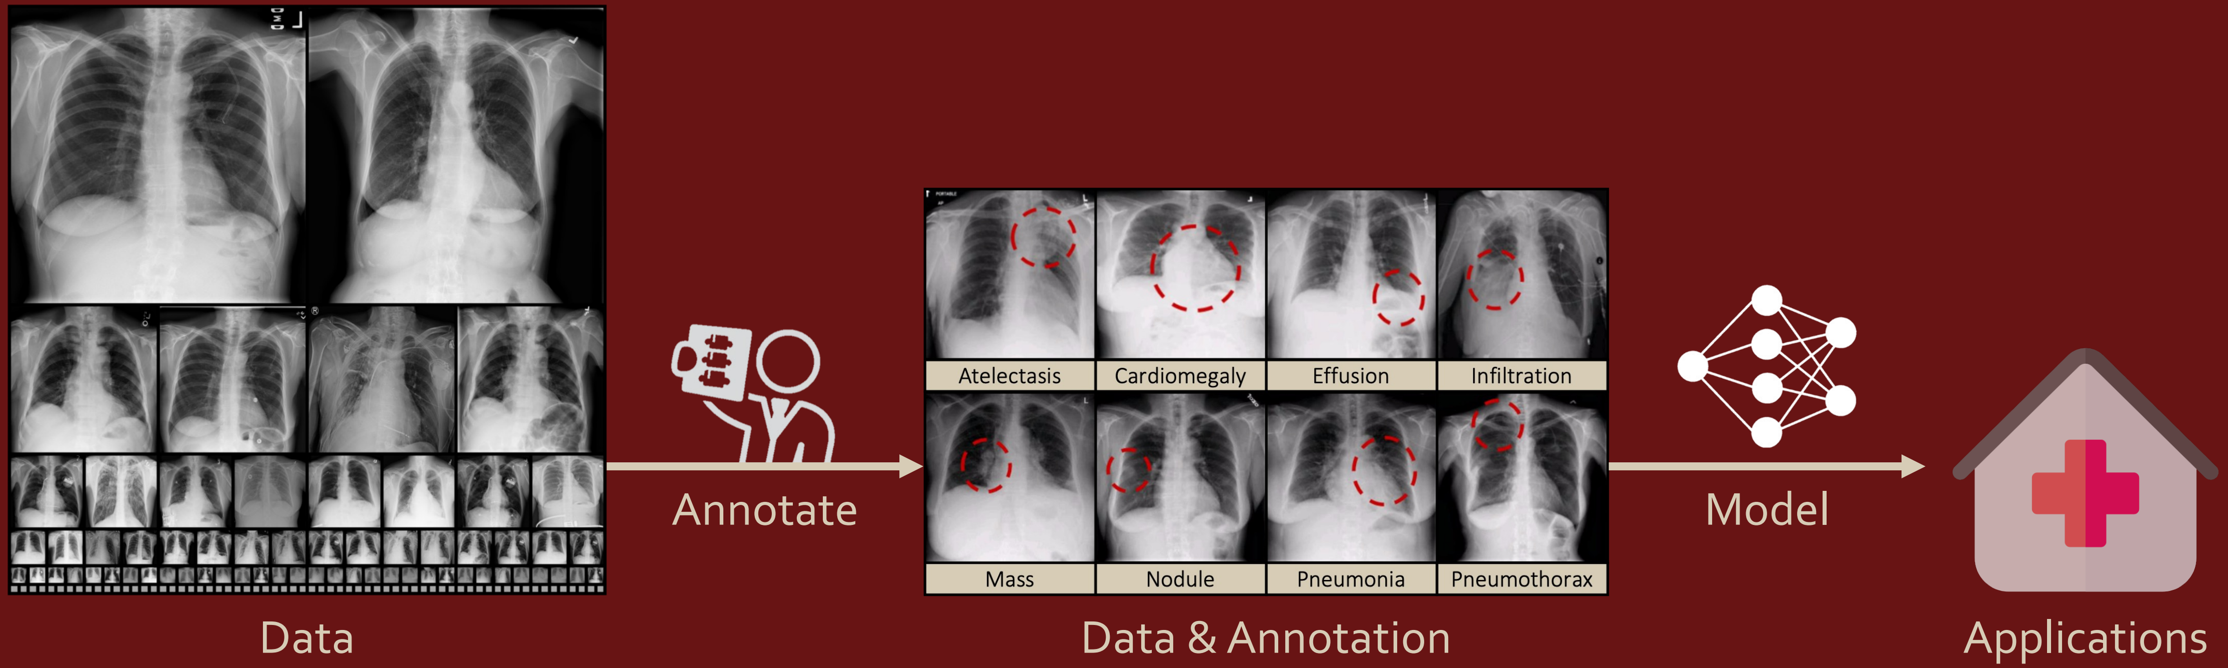
Annotate
Atelectasis
Cardiomegaly
Effusion
Infiltration
Mass
Nodule
Pneumonia
Pneumothorax
Model
Data
Data & Annotation
Applications

Goal: Reduce annotation efforts for radiologists.

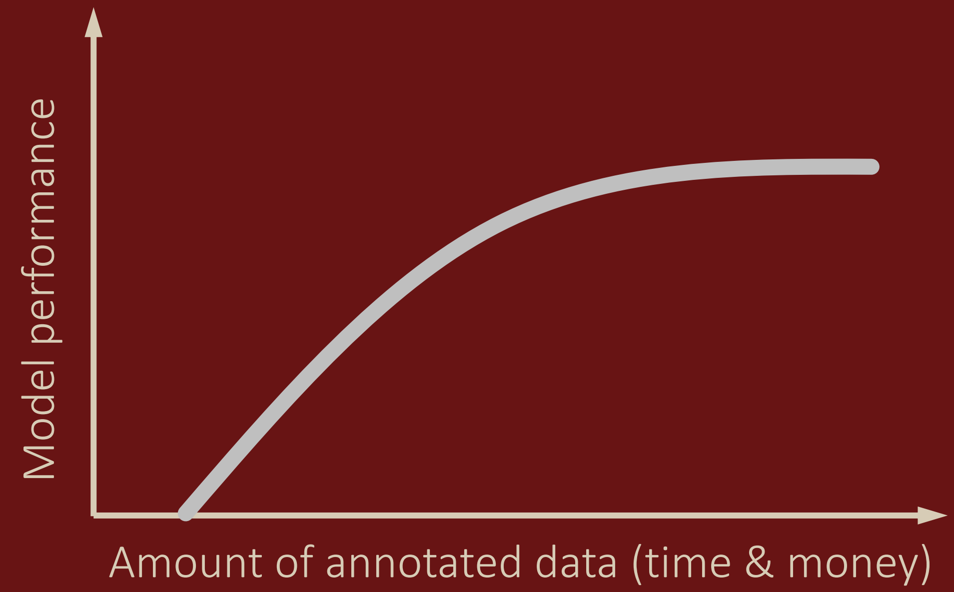

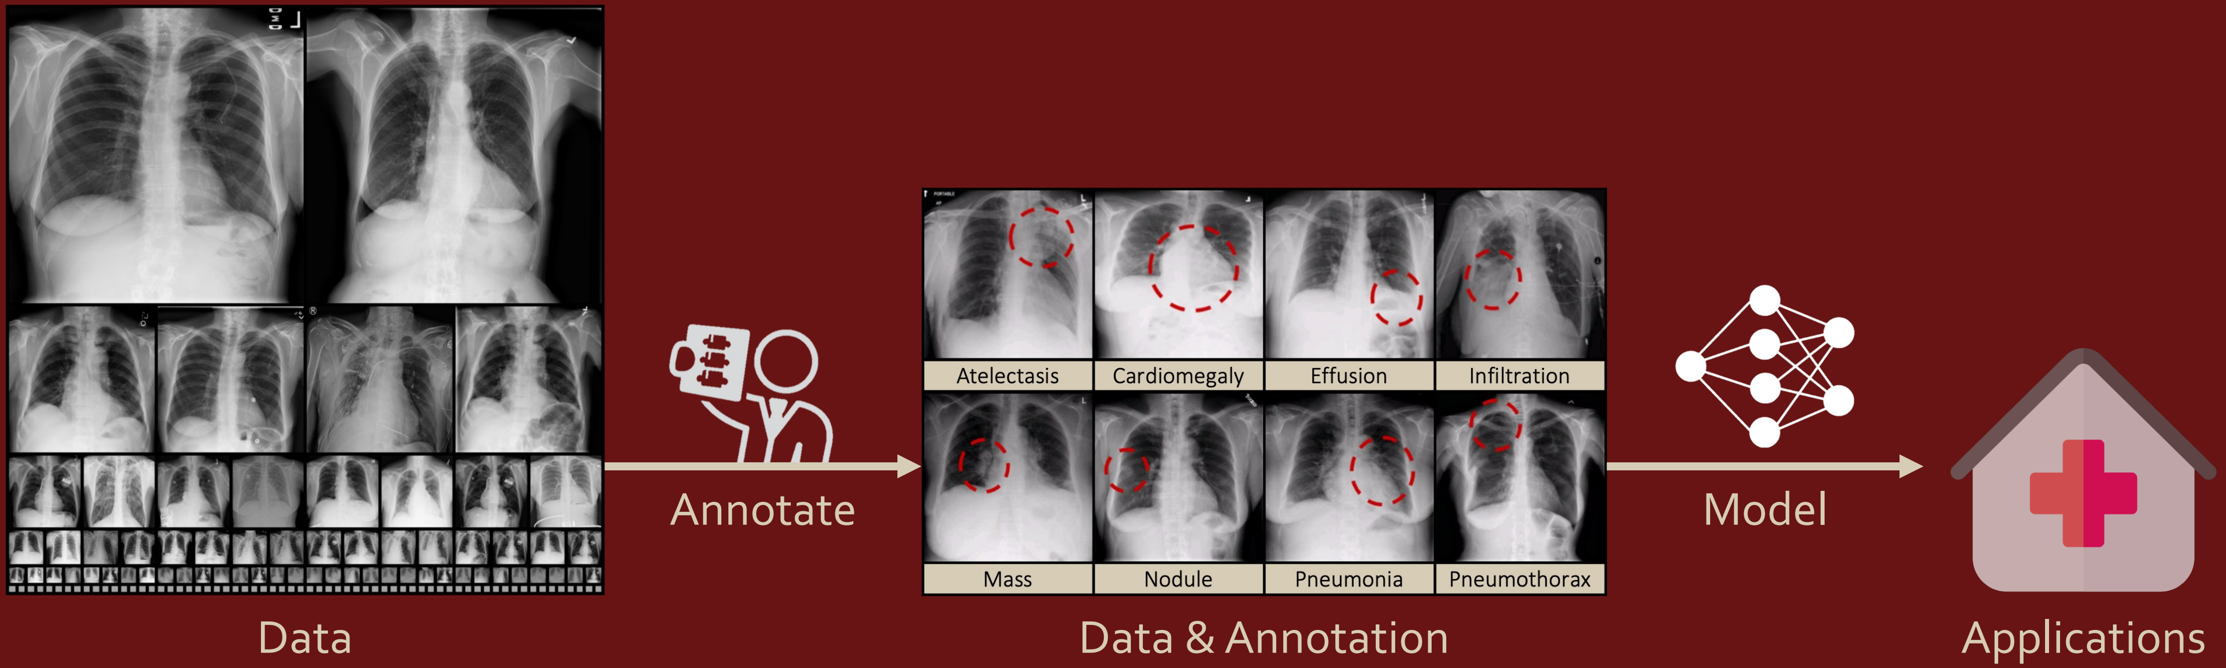

Goal: Reduce annotation efforts for radiologists.

1. Acquiring necessary annotation efficiently from human experts.

- **Active, Continual Fine-Tuning (ACFT)**
- CVPR'17, MedIA'21, MIDL'23

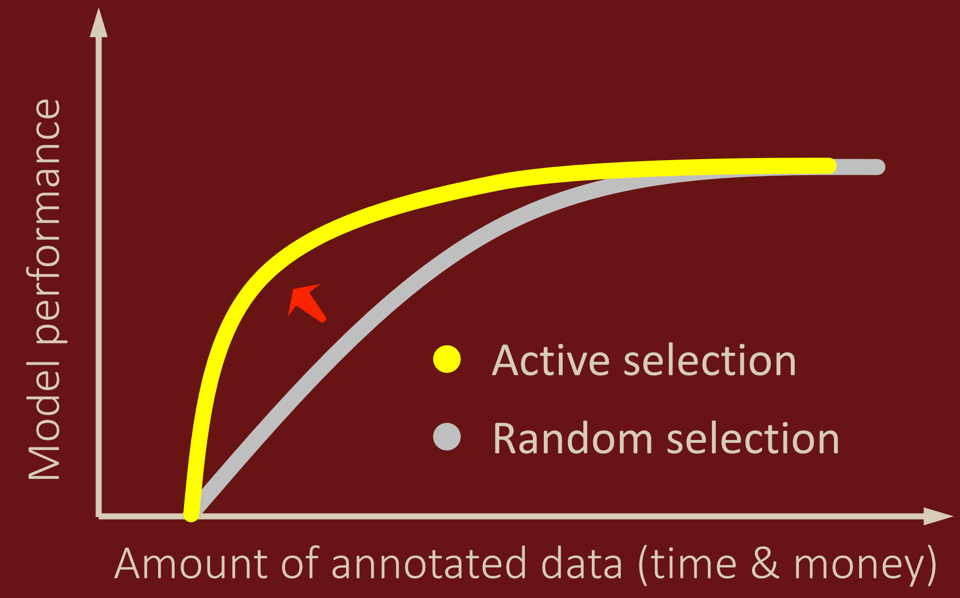

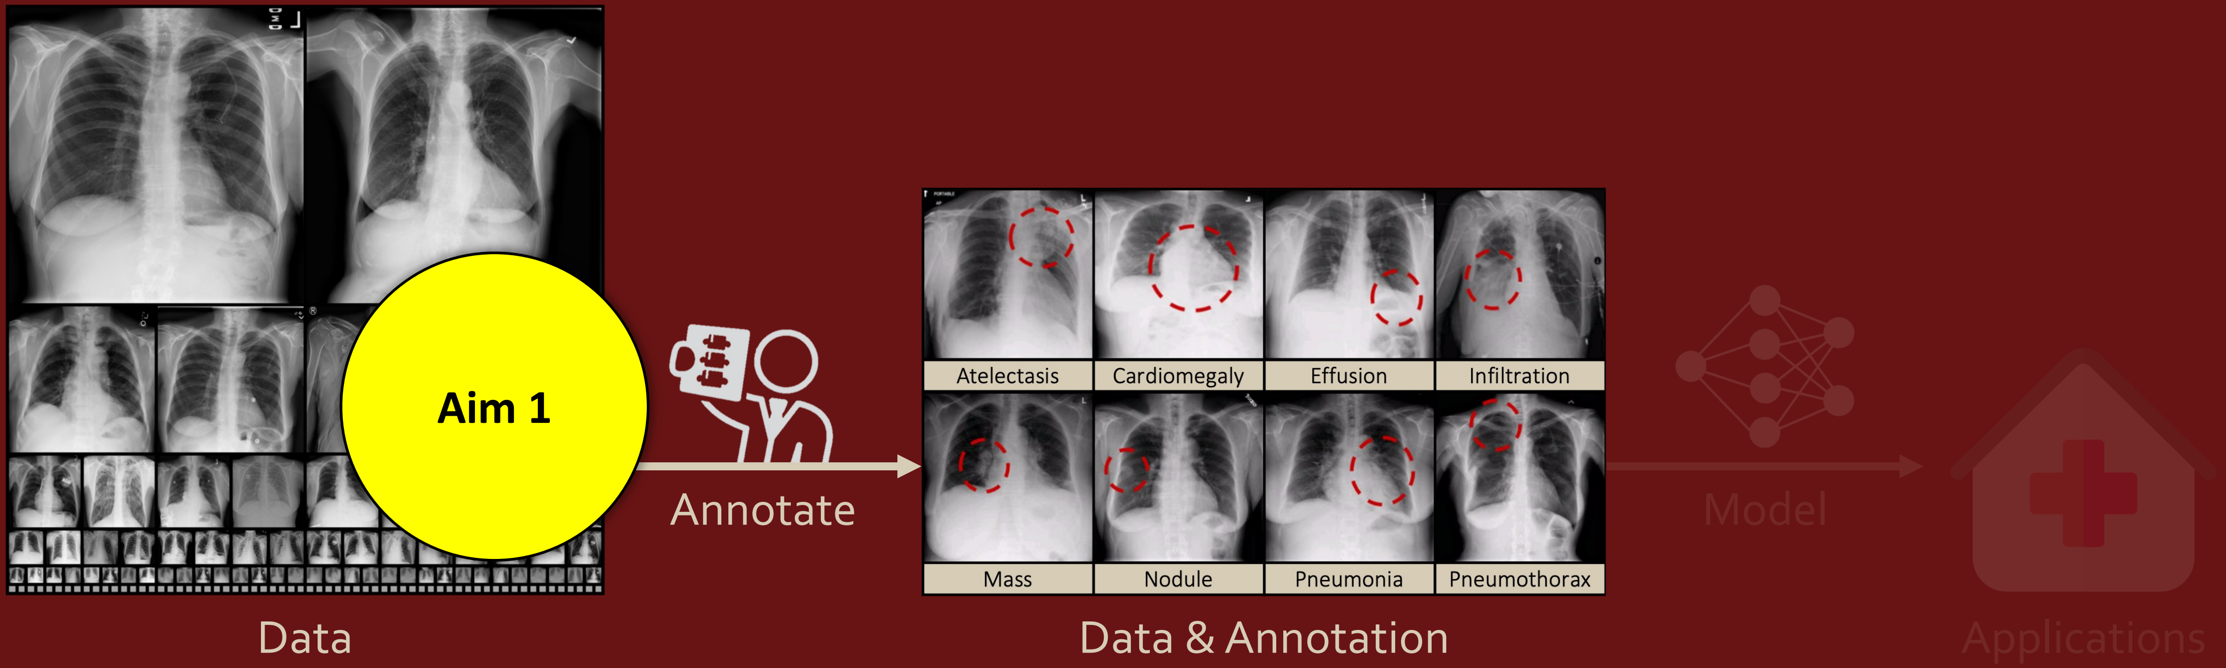

Goal: Reduce annotation efforts for radiologists.

1. Acquiring necessary annotation efficiently from human experts.

- **Active, Continual Fine-Tuning (ACFT)**
- CVPR'17, MedIA'21, MIDL'23
- Integrating uncertainty and diversity criteria
- Reducing over **80%** annotation cost

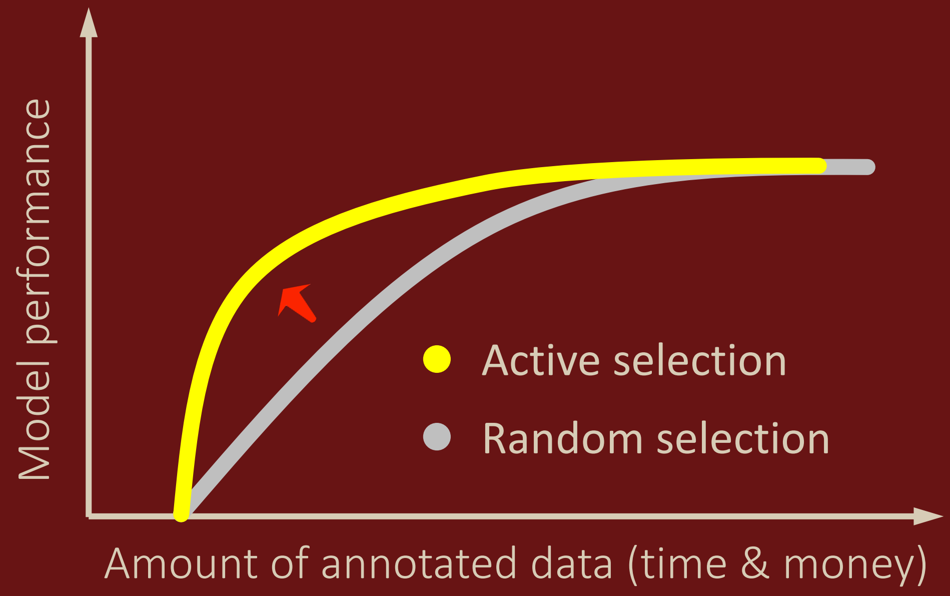

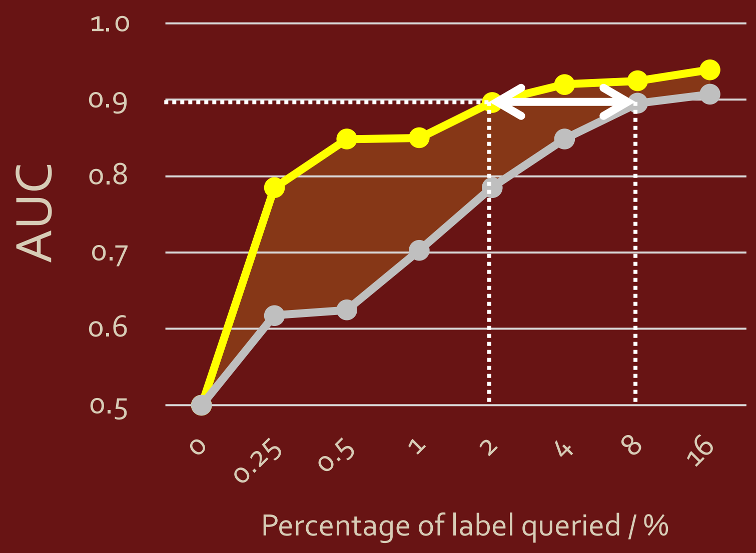

Colonoscopy Frame Classification (-81.5%)

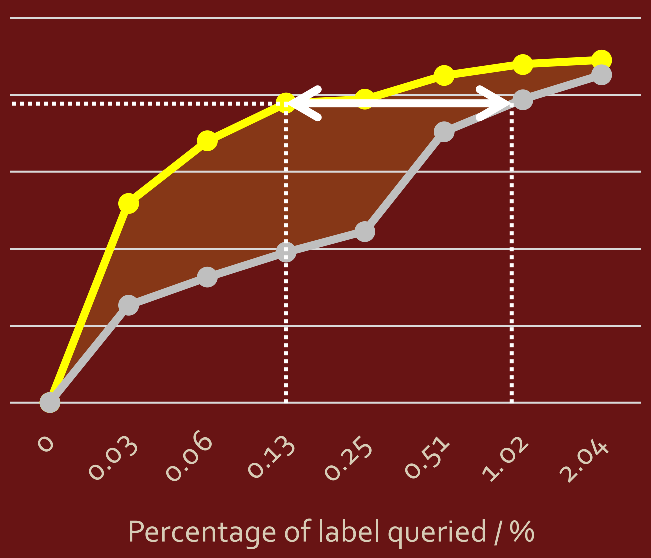

Polyp Detection (-86.3%)

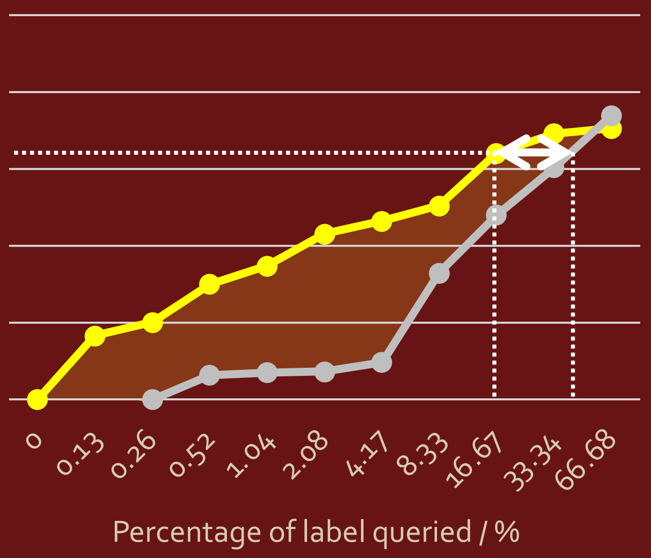

Pulmonary Embolism Detection (-80.3%)

Goal: Reduce annotation efforts for radiologists.

1. Acquiring necessary annotation efficiently from human experts.
2. Utilizing existing annotation effectively from advanced models.

- **UNet++**
- MICCAIW’18, IEEE TMI’19 (*Most Popular Articles*)

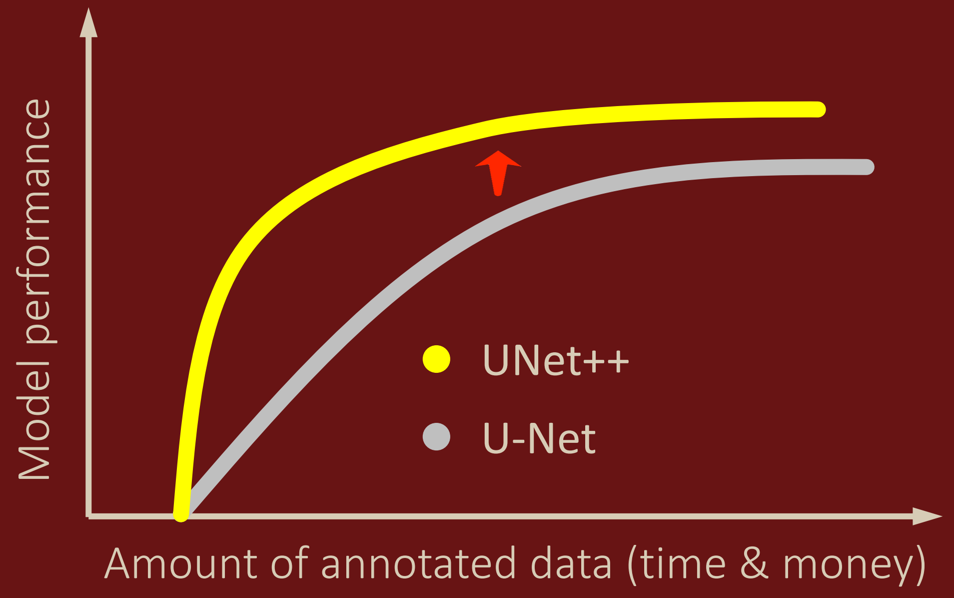

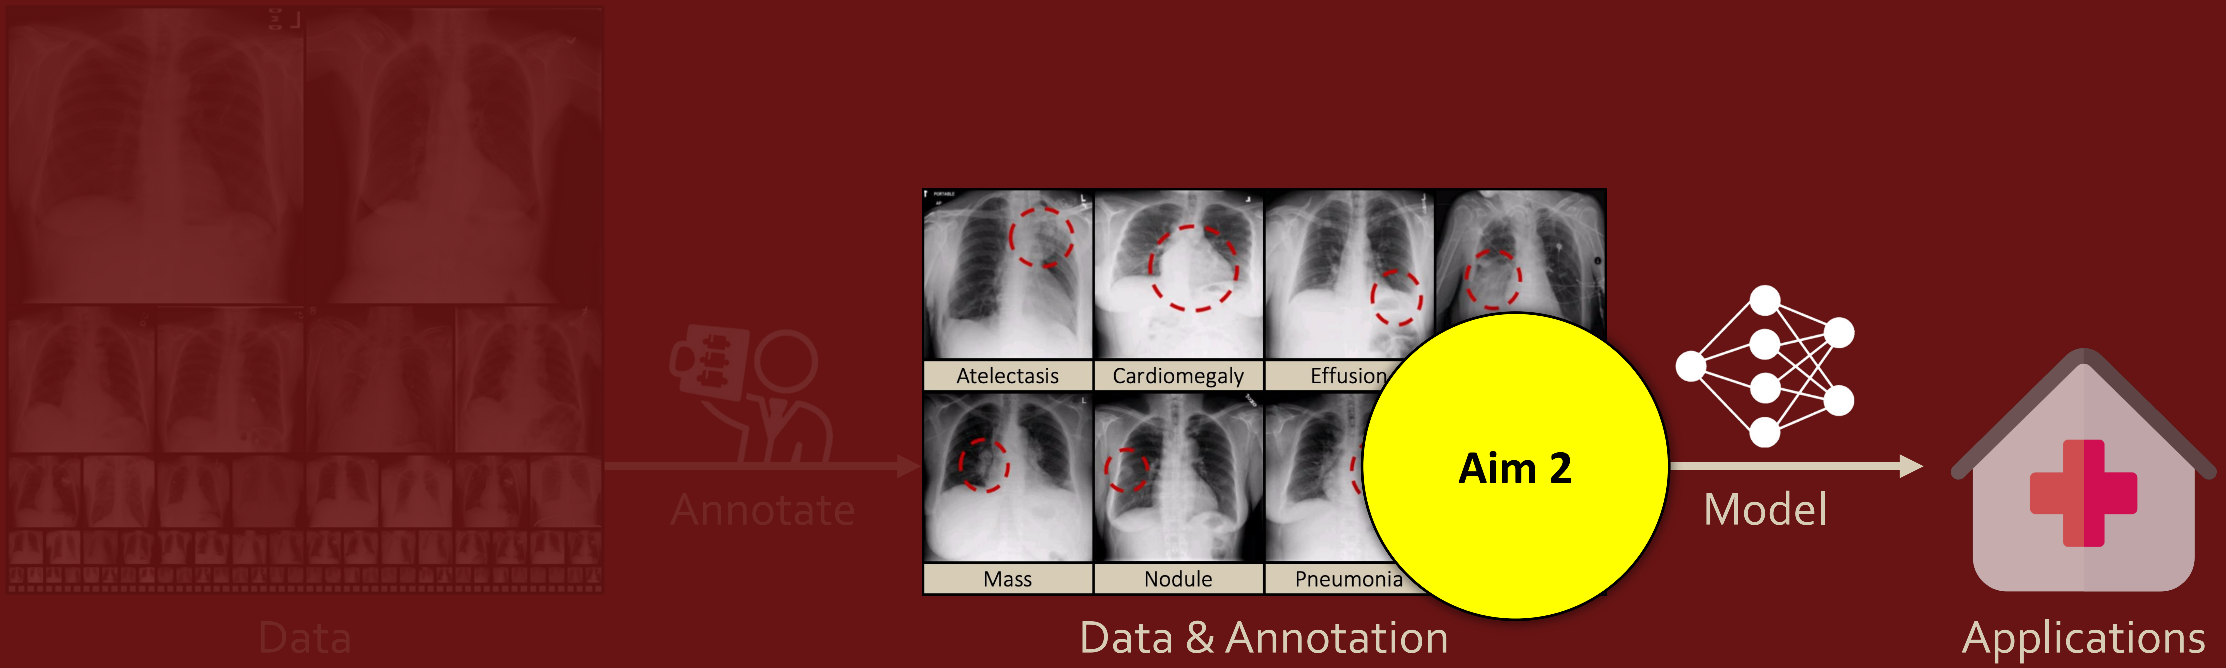

Goal: Reduce annotation efforts for radiologists.

1. Acquiring necessary annotation efficiently from human experts.
2. Utilizing existing annotation effectively from advanced models.
    - **UNet++**
    - MICCAIW’18, IEEE TMI’19 (*Most Popular Articles*)
    - Aggregating multi-scale, multi-resolution features
    - Detecting **very small tumors** without too many FPs

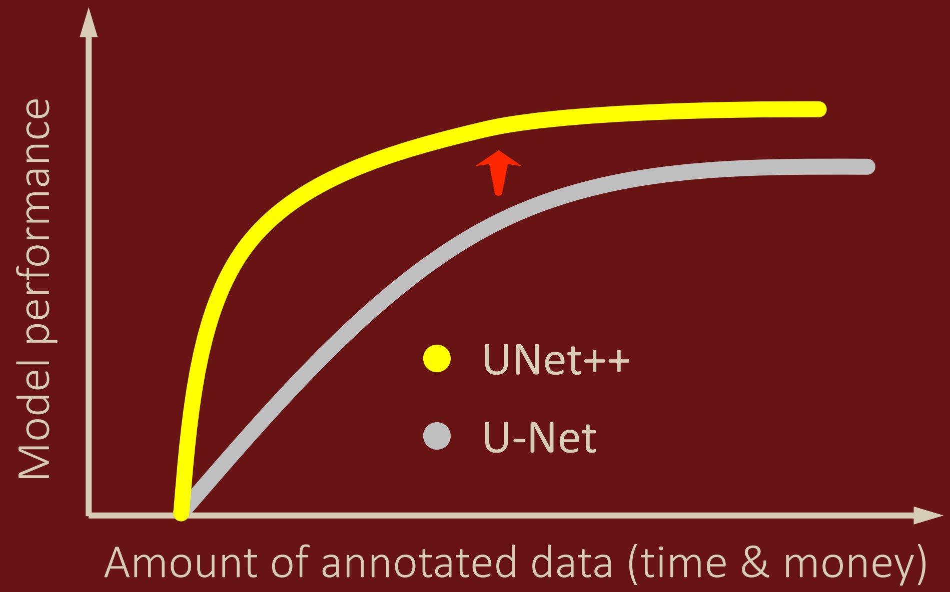

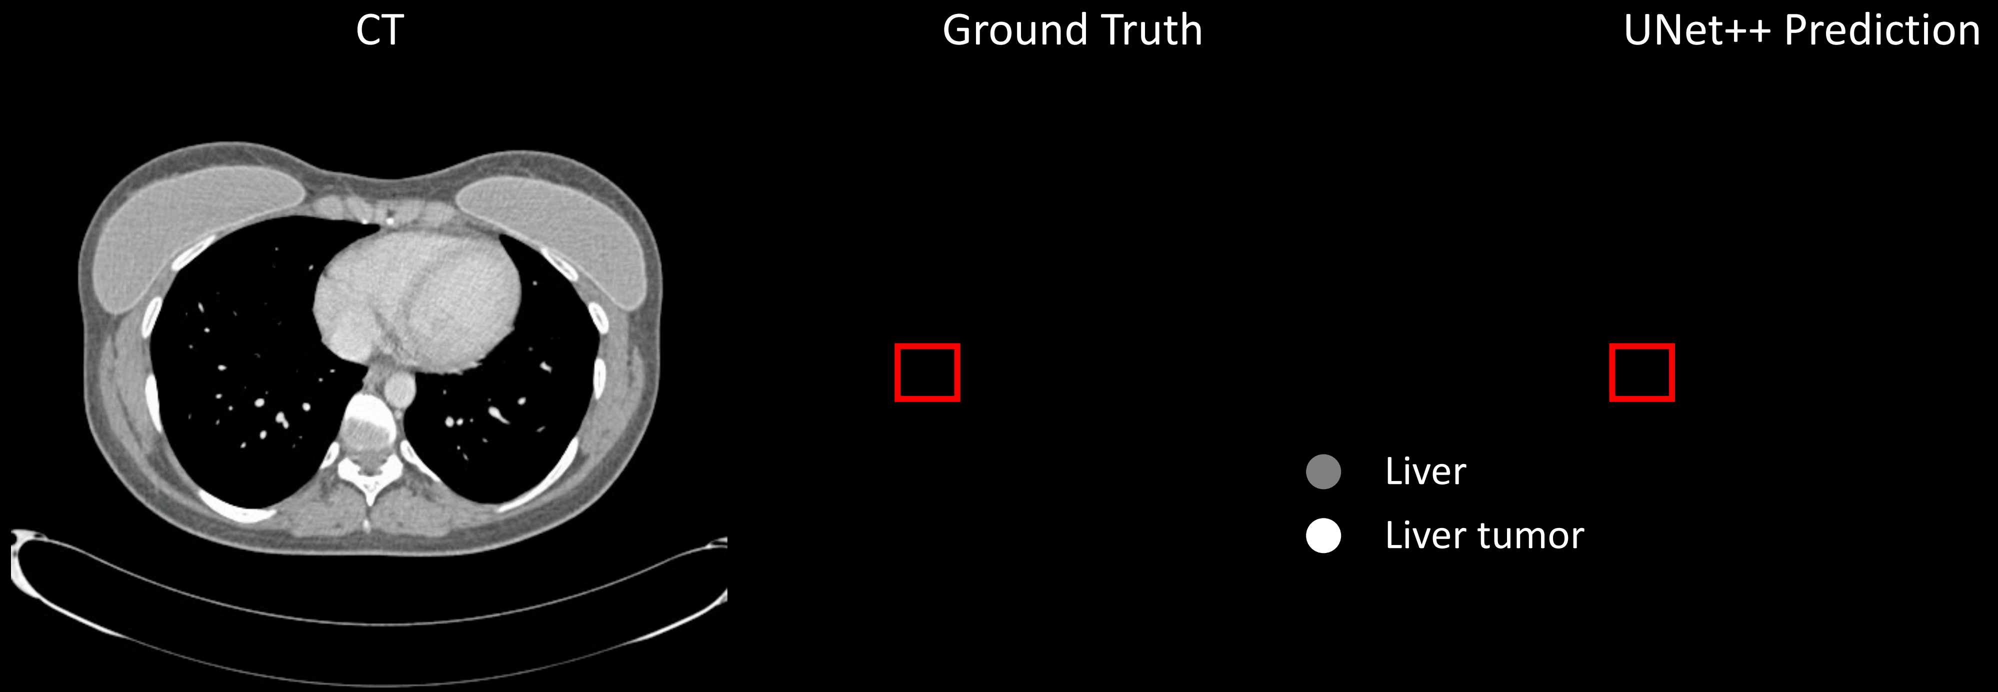

Goal: Reduce annotation efforts for radiologists.

1. Acquiring necessary annotation efficiently from human experts.
2. Utilizing existing annotation effectively from advanced models.
    - **UNet++**
    - MICCAIW’18, IEEE TMI’19 (*Most Popular Articles*)
    - Aggregating multi-scale, multi-resolution features
    - Detecting **very small tumors** without too many FPs

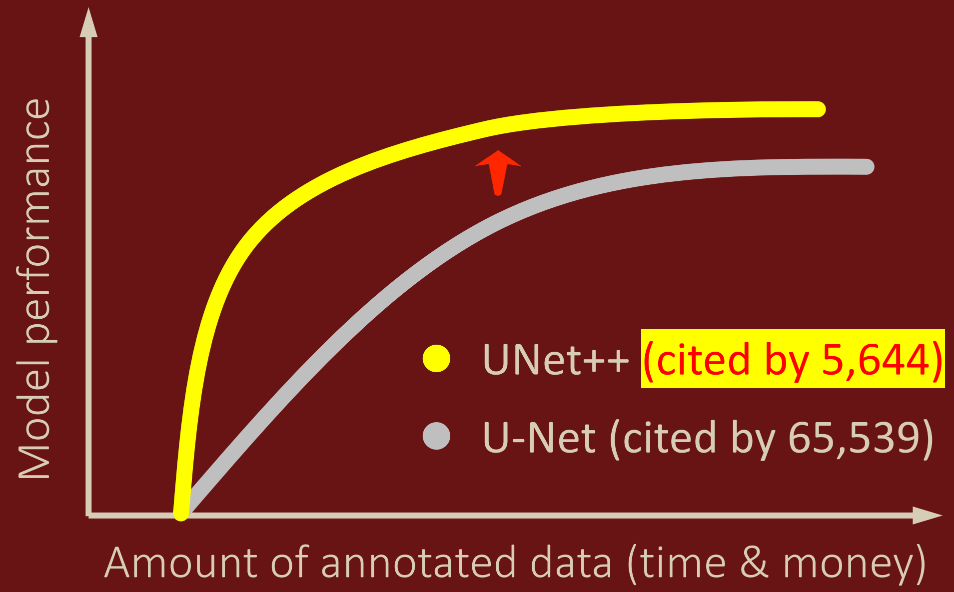

43.9% → 58.1% (U-Net → UNet++) Covid-19 segmentation (CT) [Fan et al., IEEE TMI]	**78.6% → 82.9% (U-Net → UNet++)** Fiber tracing (corneal confocal microscopy) [Mou et al., MICCAI]	**86.5% → 89.5% (U-Net → UNet++)** Spleen segmentation (MRI) [Li et al., Computers & Graphics]
86.6% → 87.2% (U-Net → UNet++) SegTHOR 2019 Challenge (CT) [Zhang et al., IEEE TMI]	**90.2% → 92.0% (U-Net → UNet++)** Optic Disc & Cup Segmentation (fundus image) [Meng et al., MICCAI]	**60.3% → 71.6% (U-Net → UNet++)** Ground-glass opacity segmentation (CT) [Zheng et al., IEEE Access]
51.2% → 58.6% (U-Net → UNet++) Esophagus segmentation (CT) [Huang et al., IEEE Access]	**63.7% → 66.3% (U-Net → UNet++)** Liver tumor segmentation (CT) [Bajpai et al., Master Thesis]	**90.7% → 91.6% (U-Net → UNet++)** Heart segmentation (MRI) [Ji et al., MICCAI]

Goal: Reduce annotation efforts for radiologists.

1. Acquiring necessary annotation efficiently from human experts.
2. Utilizing existing annotation effectively from advanced models.
3. Extracting generic knowledge directly from unannotated images.
   - **Models Genesis**
   - MICCAI'19 (*Young Scientist Award*), MIA (*Best Paper Award*)

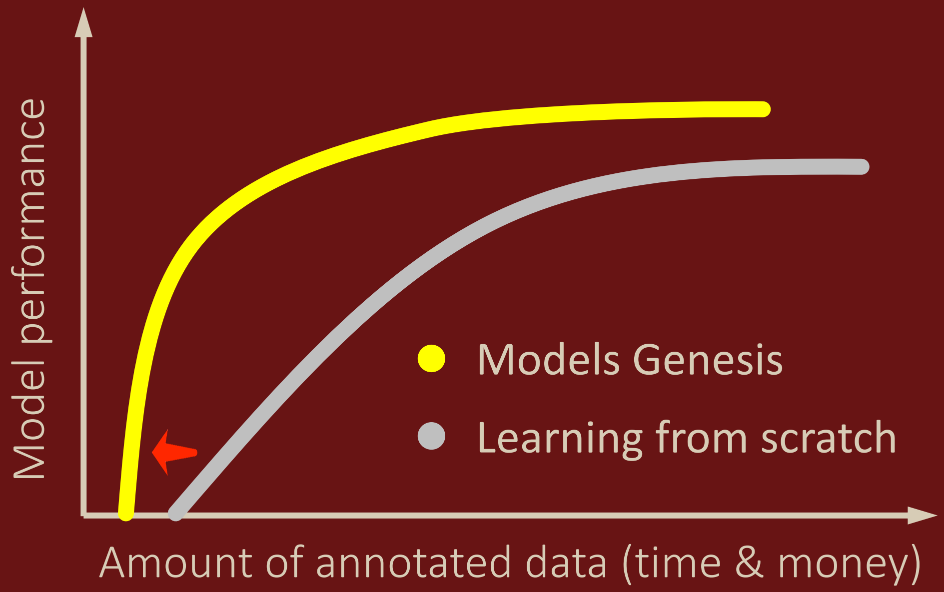

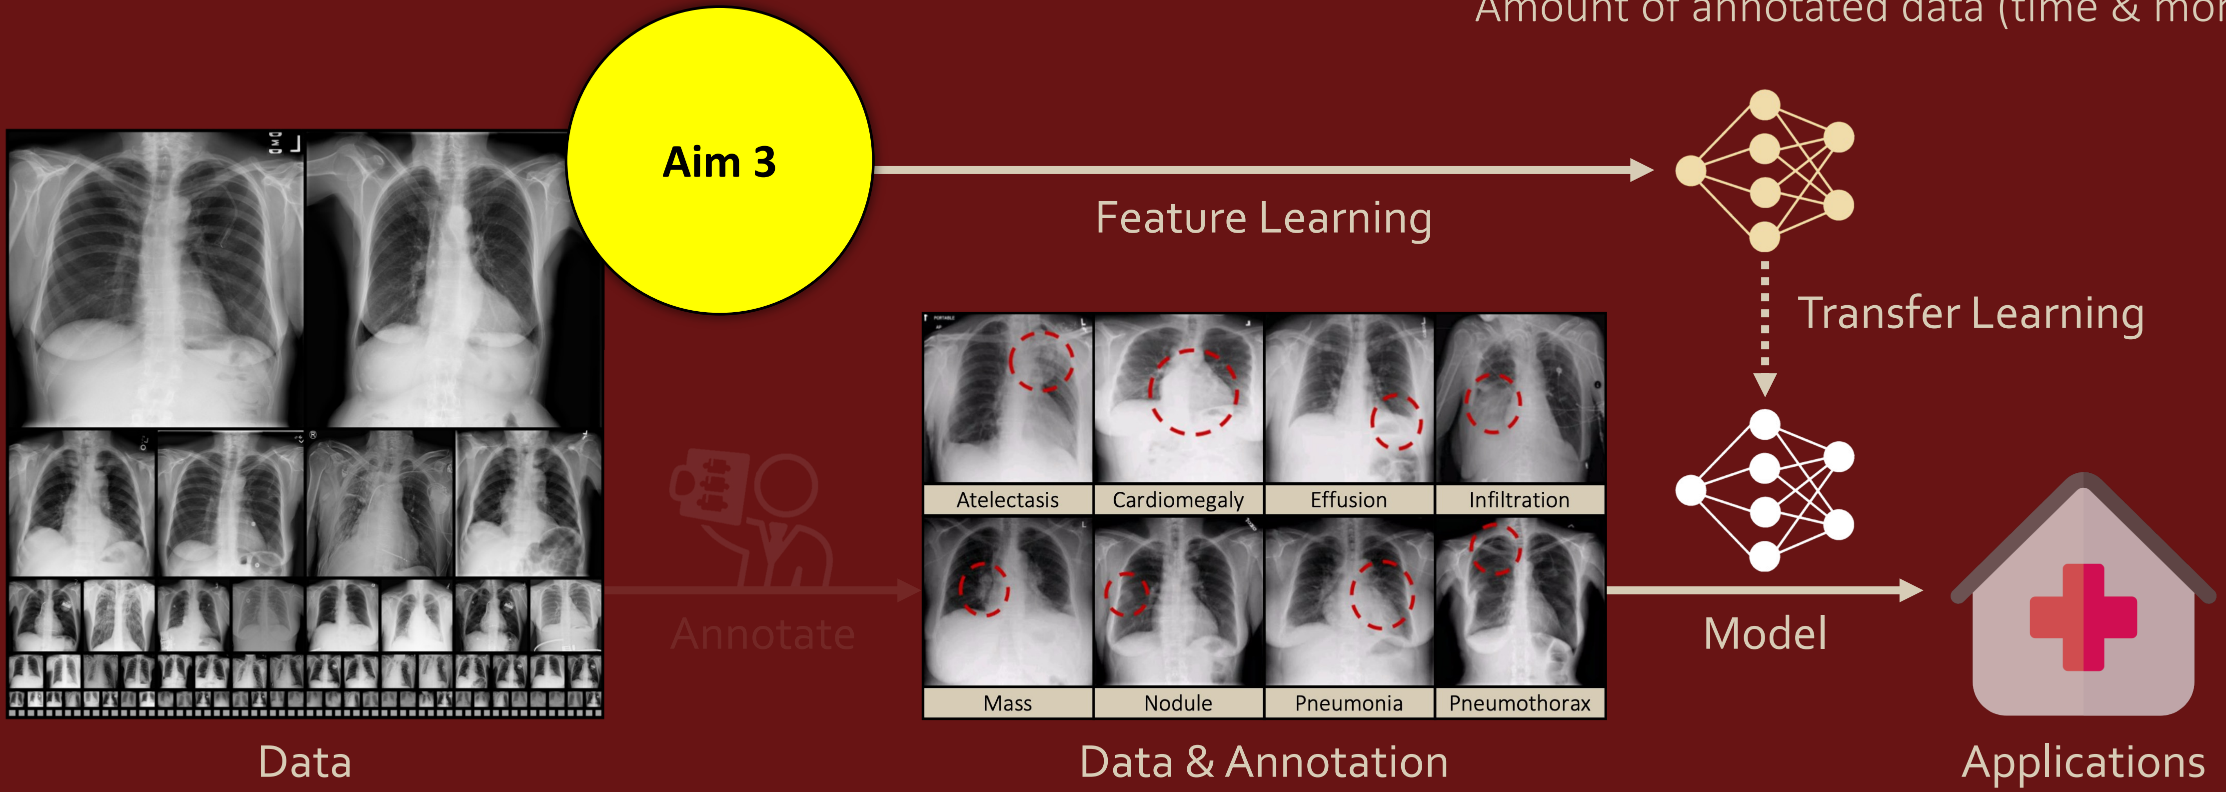

Goal: Reduce annotation efforts for radiologists.

1. Acquiring necessary annotation efficiently from human experts.
2. Utilizing existing annotation effectively from advanced models.
3. Extracting generic knowledge directly from unannotated images.
   - **Models Genesis**
   - MICCAI'19 (*Young Scientist Award*), MIA (*Best Paper Award*)
   - The **First** publicly available 3D pre-trained model
   - A demonstration of **masked image modeling** in medicine

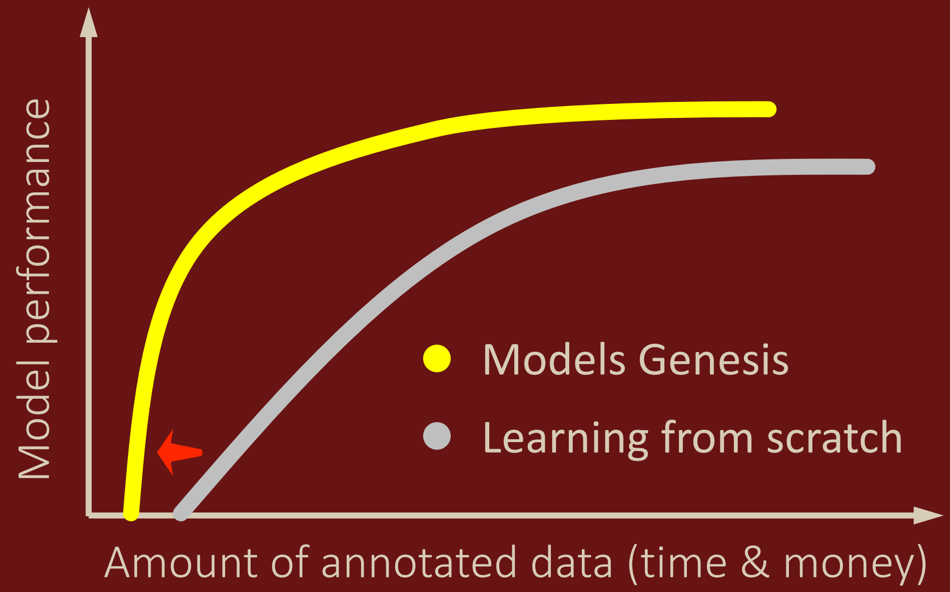

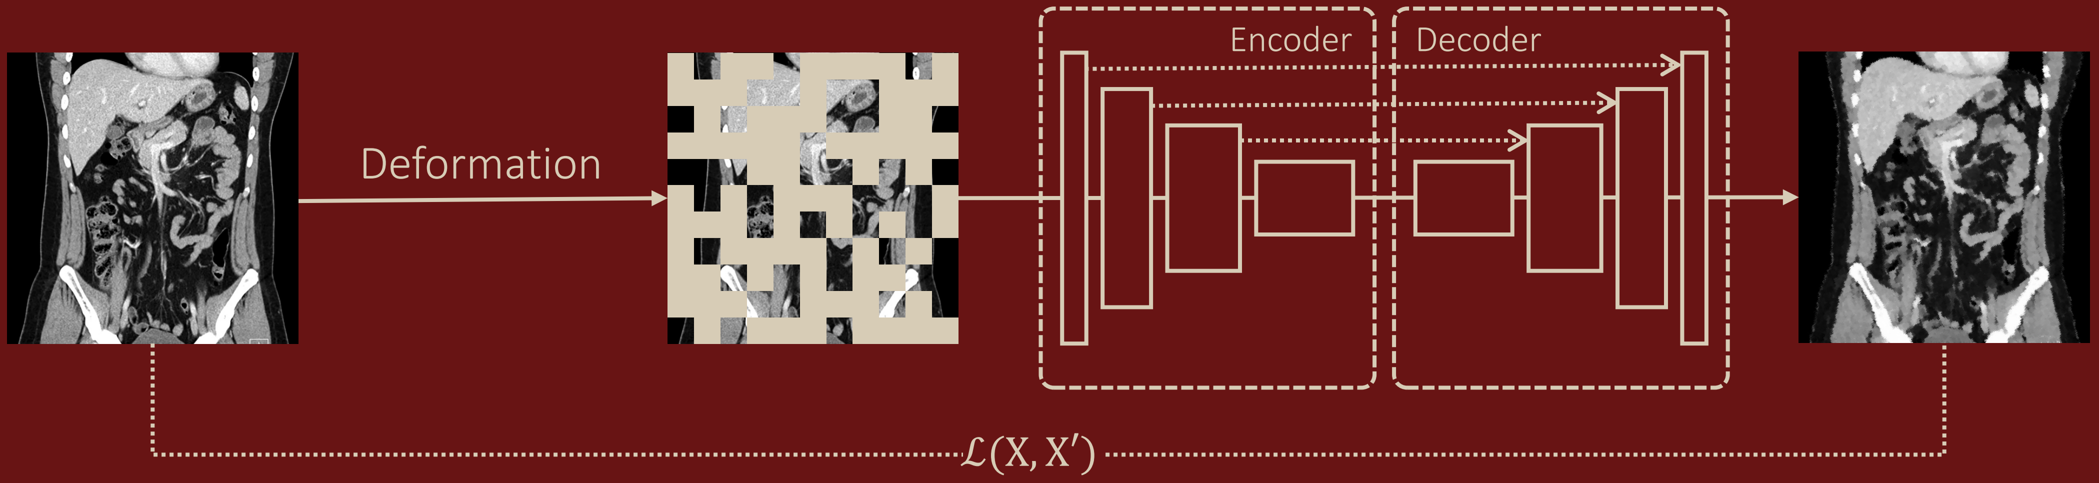

Goal: Reduce annotation efforts for radiologists.

1. Acquiring necessary annotation efficiently from human experts.
2. Utilizing existing annotation effectively from advanced models.
3. Extracting generic knowledge directly from unannotated images.

——PhD dissertation——

AMIA INFORMATICS PROFESSIONALS. LEADING THE WAY. **Doctoral Dissertation Award Winner**

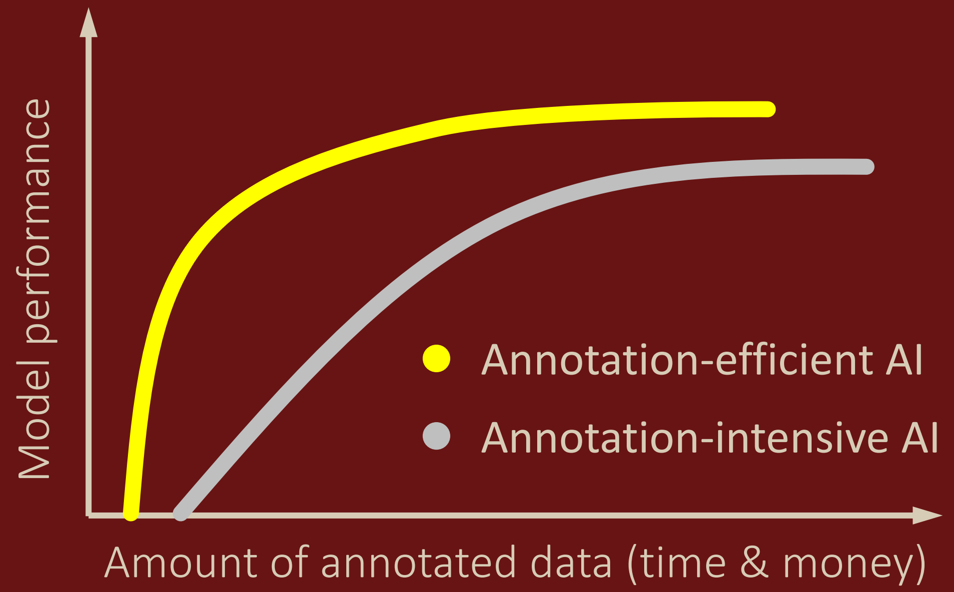

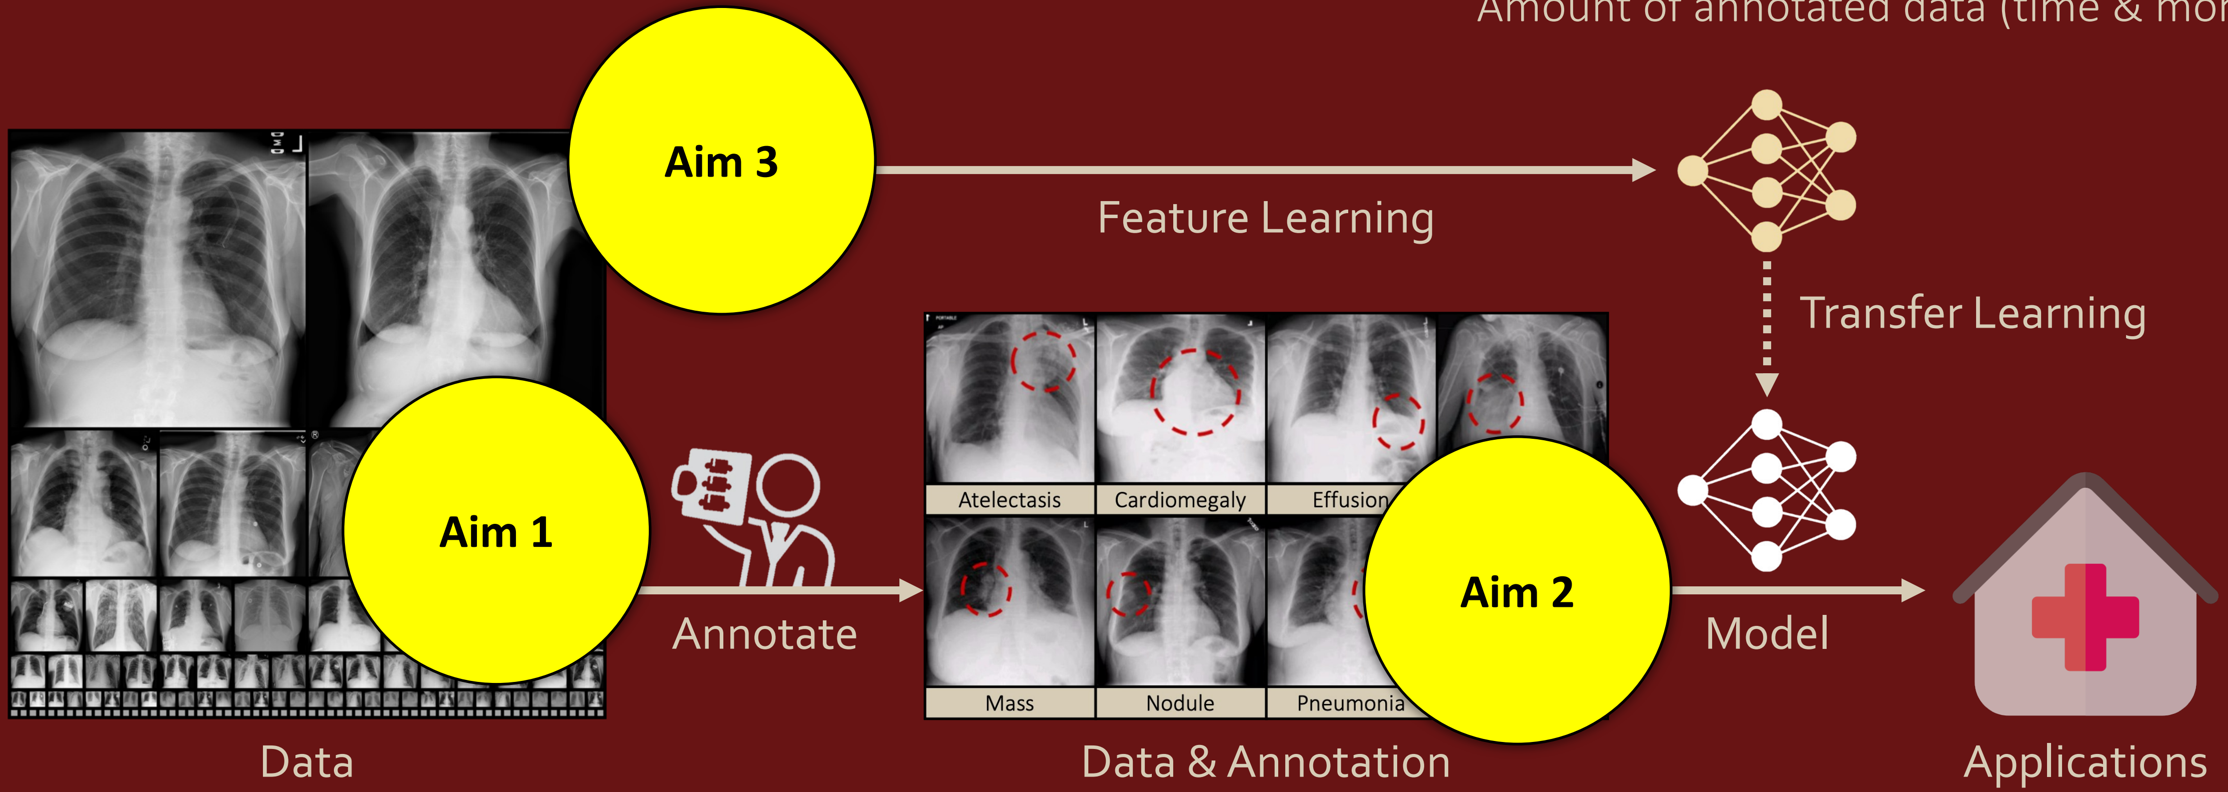

Goal: Reduce annotation efforts for radiologists.

1. Acquiring necessary annotation efficiently from human experts.
2. Utilizing existing annotation effectively from advanced models.
3. Extracting generic knowledge directly from unannotated images.

——PhD dissertation——

4. Generating annotation automatically from **tumor synthesis**.

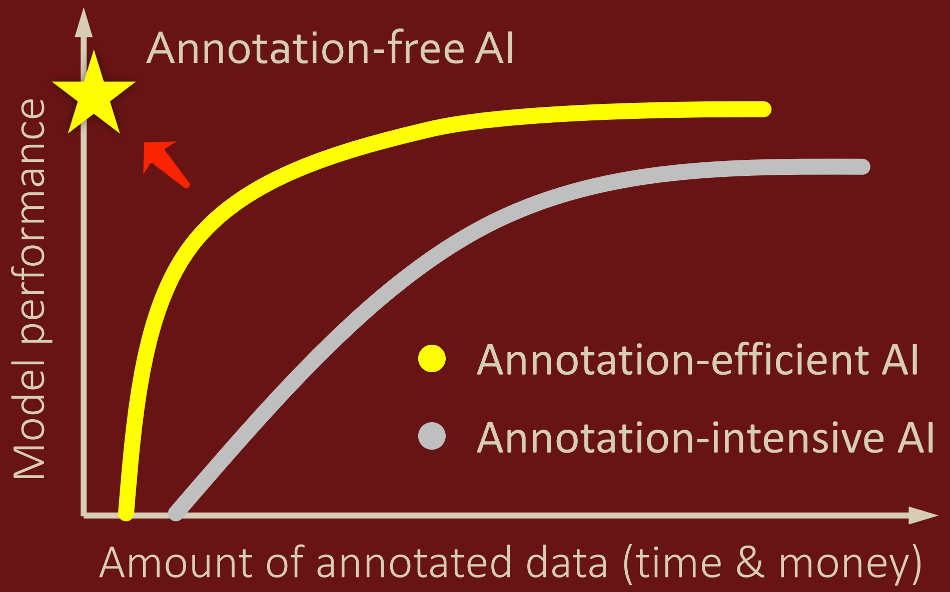

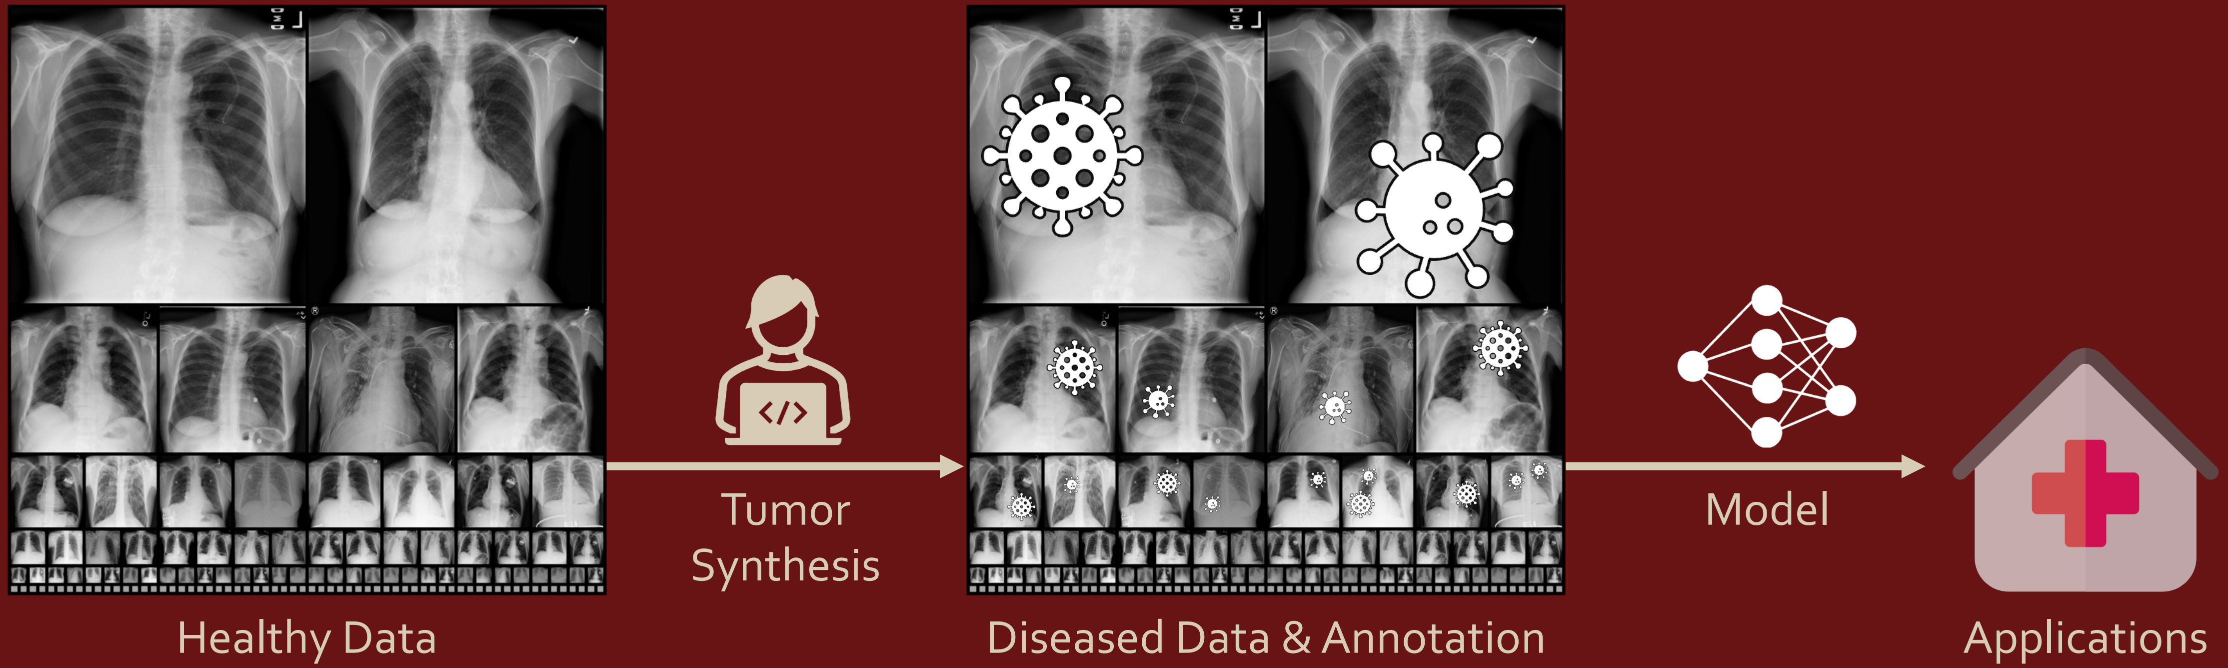

Medical professionals cannot tell which are real and which are synthetic tumors

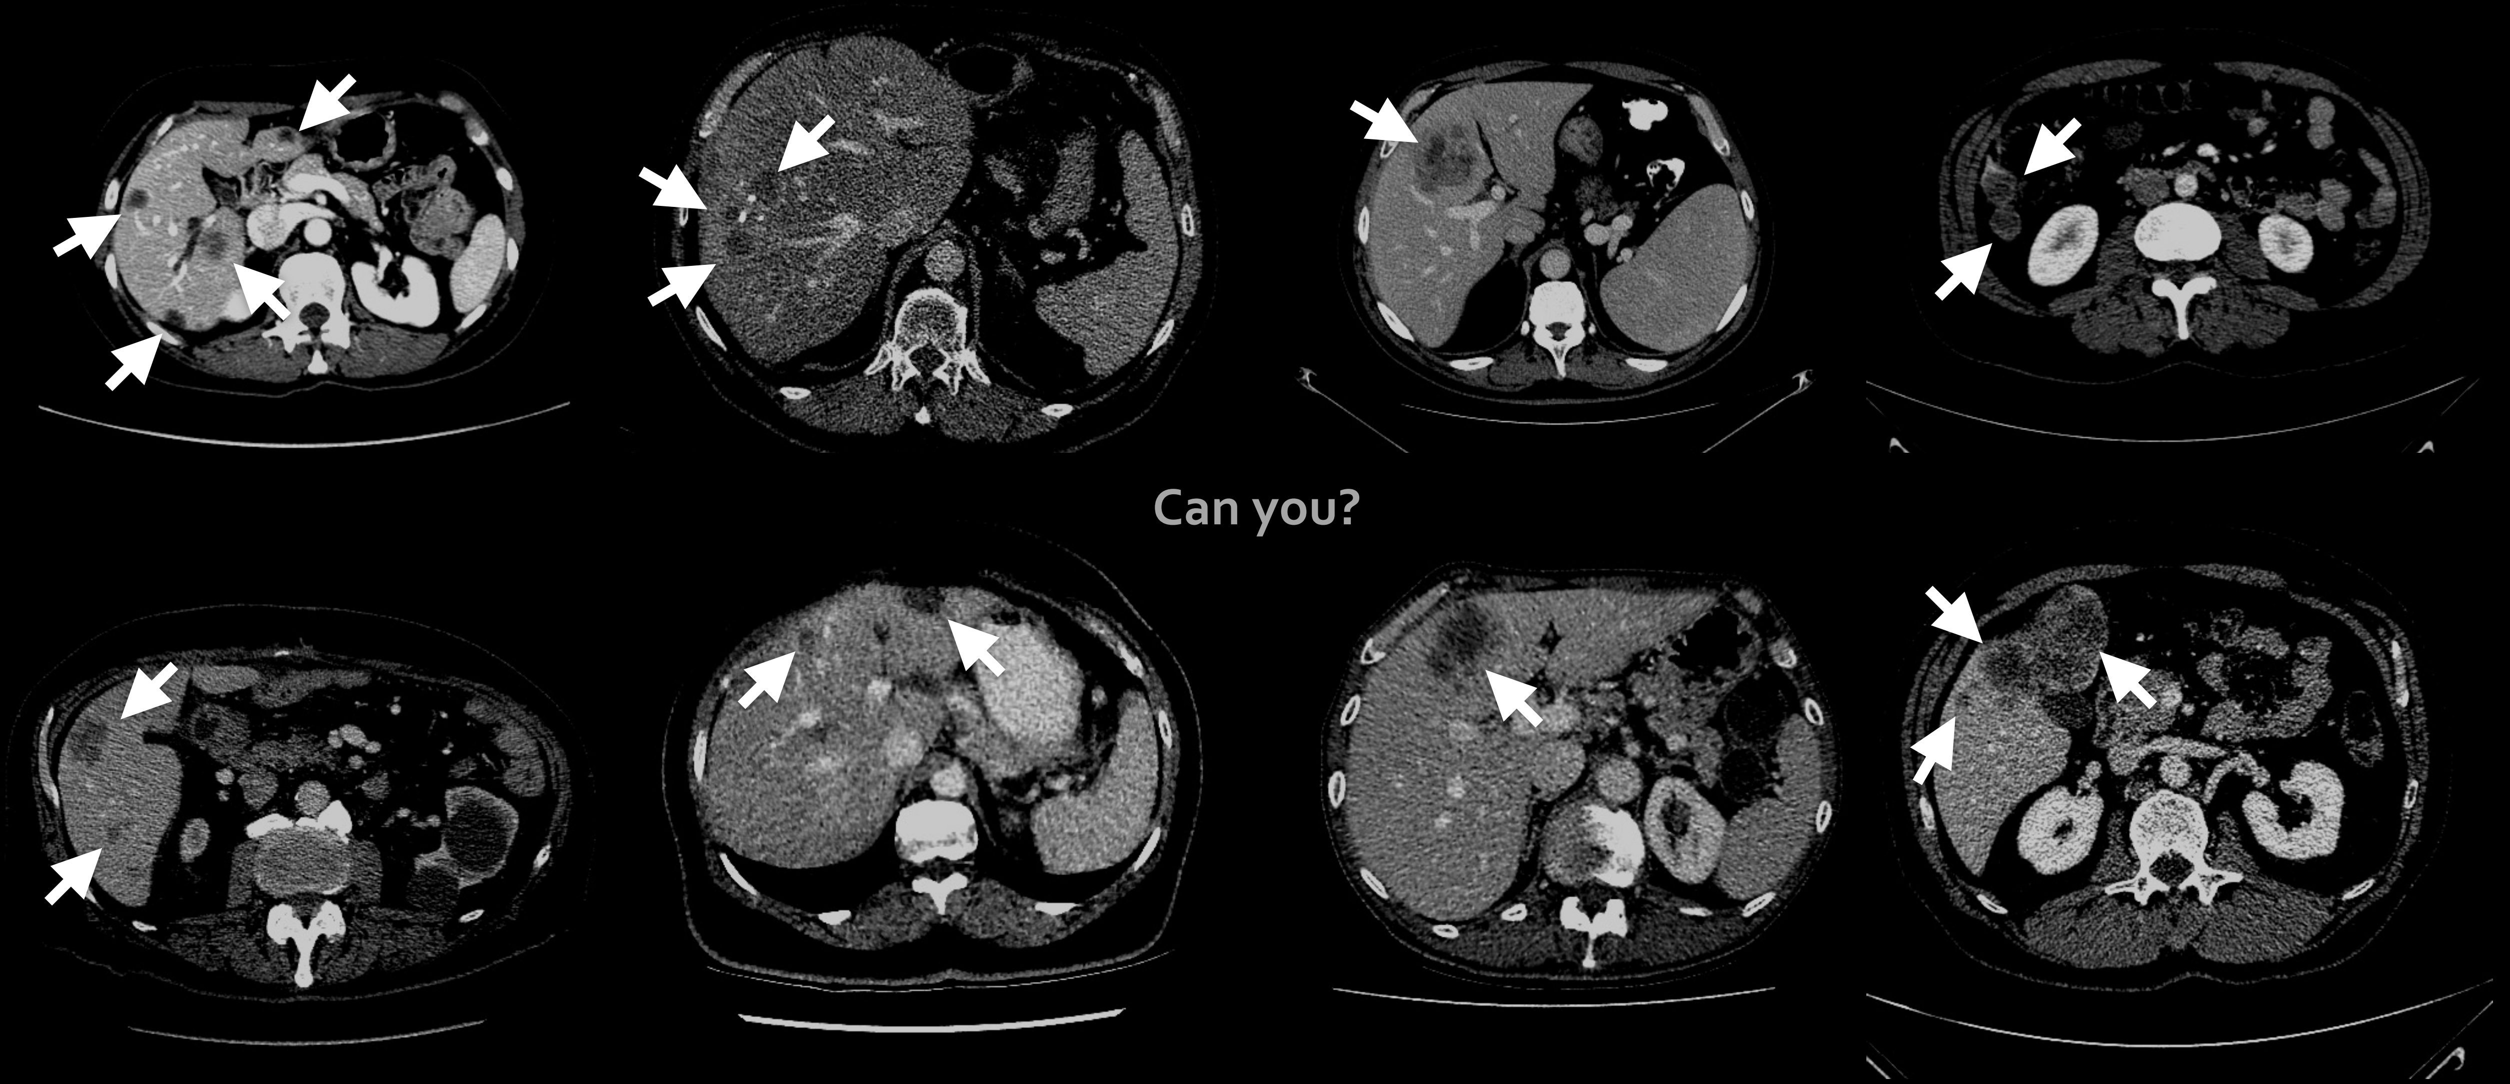

Can you?

1. Hu, Qixin, Yixiong Chen, Junfei Xiao, Shuwen Sun, Jieneng Chen, Alan Yuille, and Zongwei Zhou*. "Label-Free Liver Tumor Segmentation." CVPR-2023.

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

CT

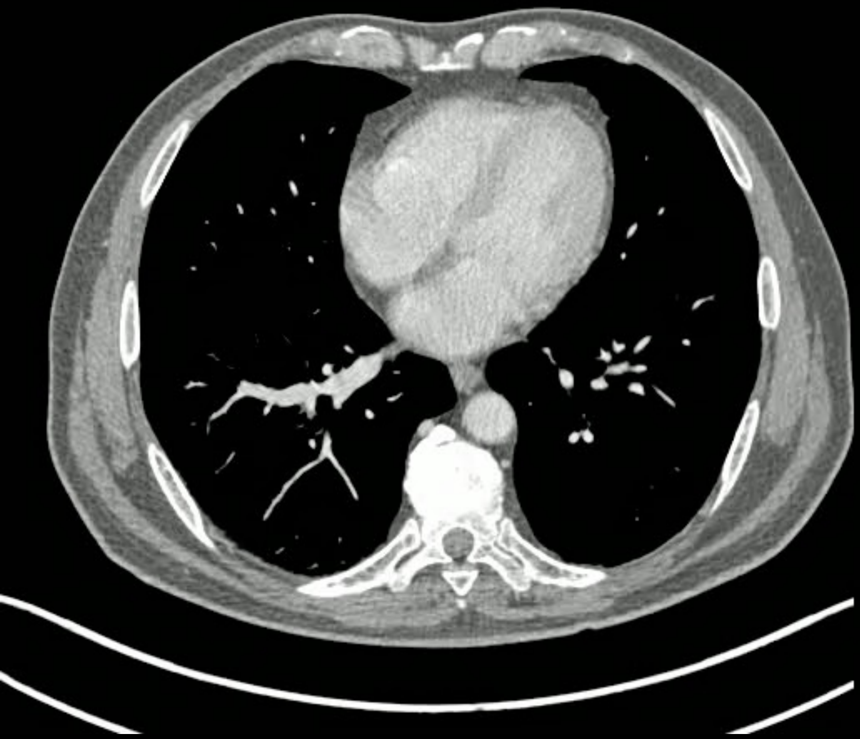

AI prediction
trained on real tumors
with per-voxel annotation
DSC = 58% [52% - 63%]

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

● Liver

● Liver tumor

1. Hu, Qixin, Yixiong Chen, Junfei Xiao, Shuwen Sun, Jieneng Chen, Alan Yuille, and Zongwei Zhou*. "Label-Free Liver Tumor Segmentation." CVPR-2023.

[Qualitative] Generating enormous small tumors for training AI models

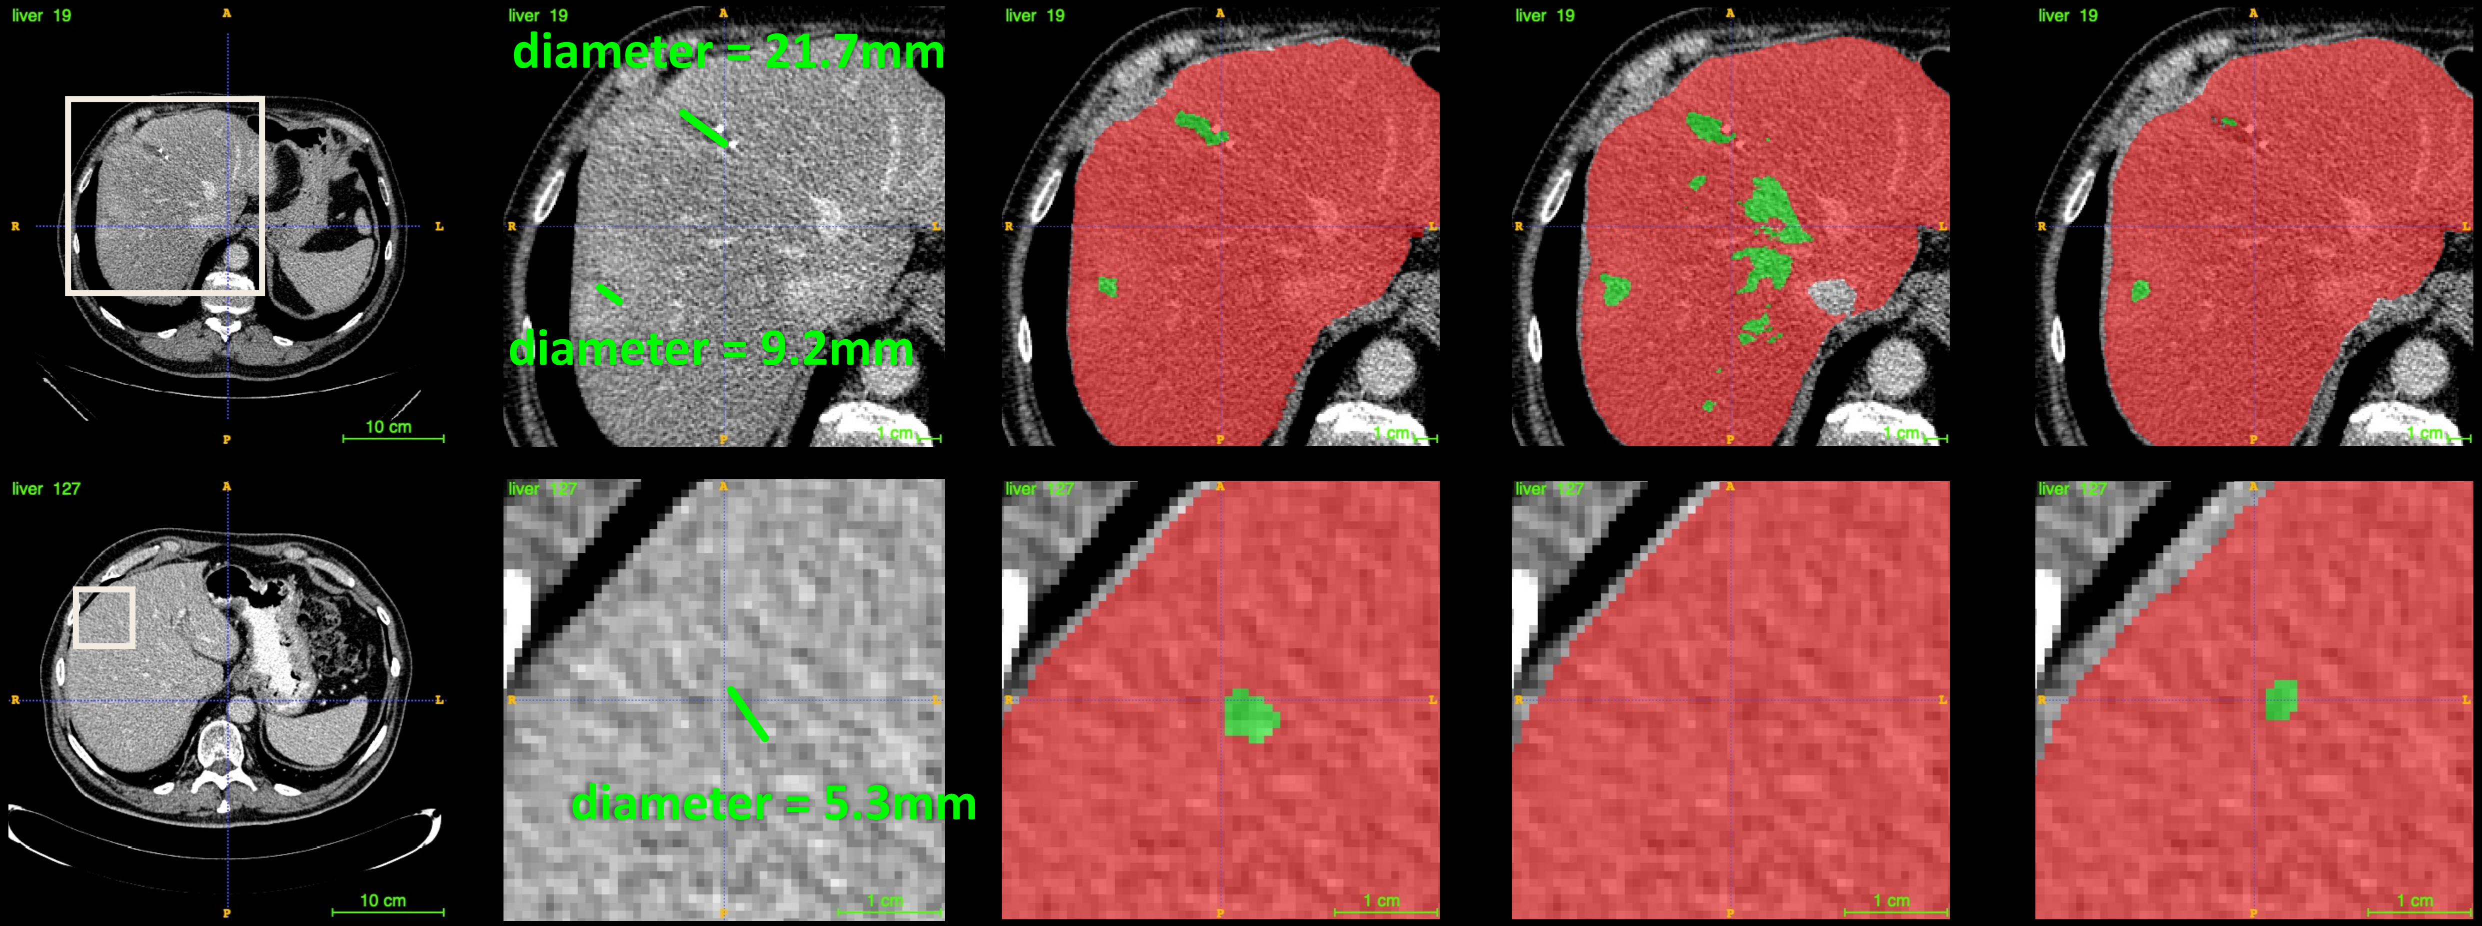

CT scan | zoom-in | ground truth | AI trained on real tumors | AI trained on synthetic tumors

1. Hu, Qixin, Yixiong Chen, Junfei Xiao, Shuwen Sun, Jieneng Chen, Alan Yuille, and Zongwei Zhou*. "Label-Free Liver Tumor Segmentation." CVPR-2023.

[Quantitative] Generating enormous small tumors for training AI models

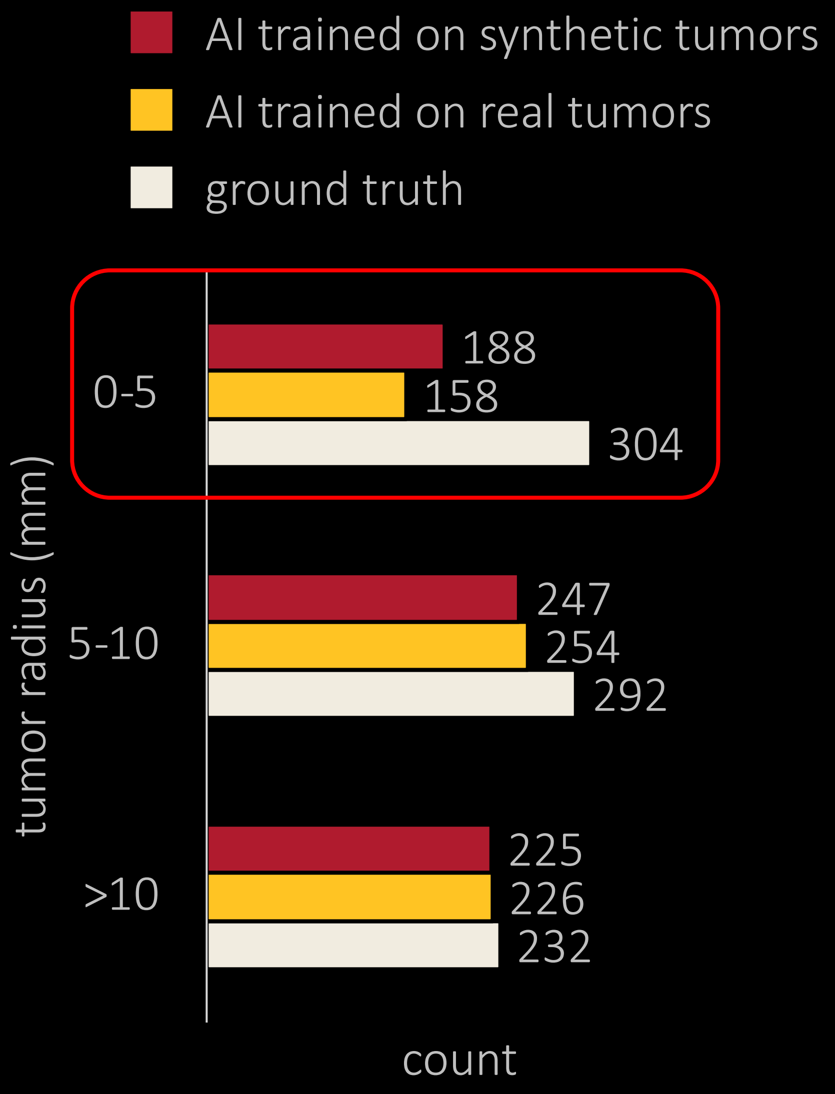

Observation: Compared with real tumors, AI trained on synthetic tumors improves Sensitivity from 52% to 62% for detecting small tumors (0-5mm).

- 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%

1. Hu, Qixin, Yixiong Chen, Junfei Xiao, Shuwen Sun, Jieneng Chen, Alan Yuille, and Zongwei Zhou*. "Label-Free Liver Tumor Segmentation." CVPR-2023.

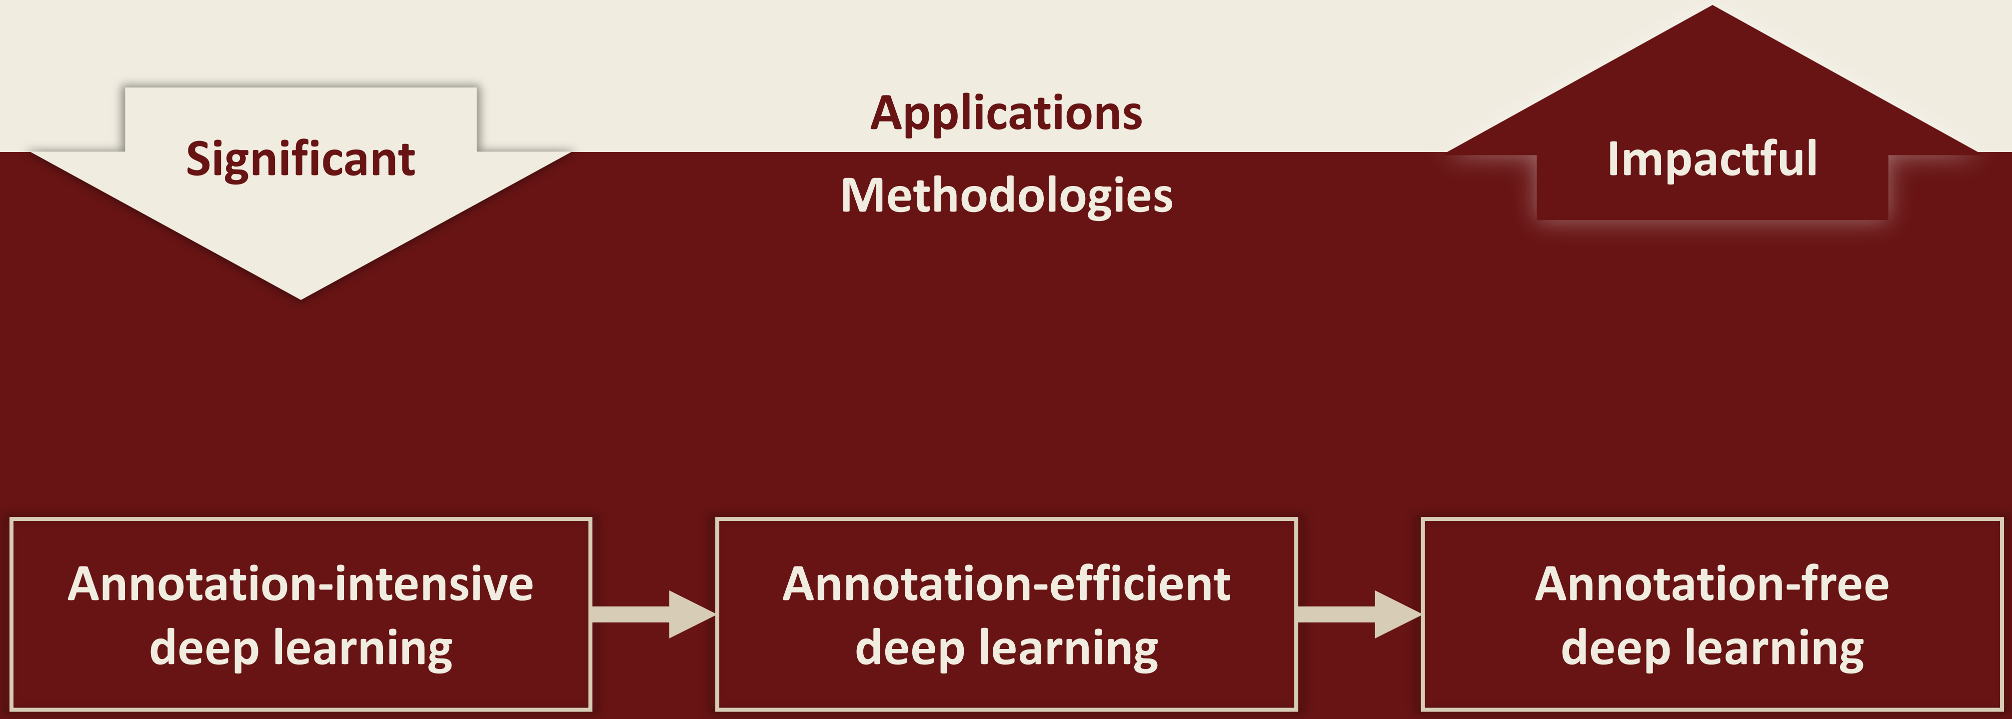
Significant
Applications
Methodologies
Impactful
Annotation-intensive deep learning
Annotation-efficient deep learning
Annotation-free deep learning

Chapter II (2020-present)

Applications: Developing 3D Maps of Whole Body

Significant

Applications

Methodologies

Impactful

Image-guided surgery	Quantitative analysis of disease progression	Earlier detection of cancer	Radiomics and predictive analytics	Neuroimaging

Edward H. Shortliffe, MD, PhD

Shortliffe Home
Professional Interests
Personal Statement
Personal Interests
C.V. & Biosketch
Textbook: Intelligent Systems in Medicine and Health (2022)
Textbook: Biomedical Informatics (5th edition, 2021)
Rule-Based Expert Systems: MYCIN (1984)
Readings in Medical Artificial Intelligence (1984)
Computer-Based Medical Consultations:

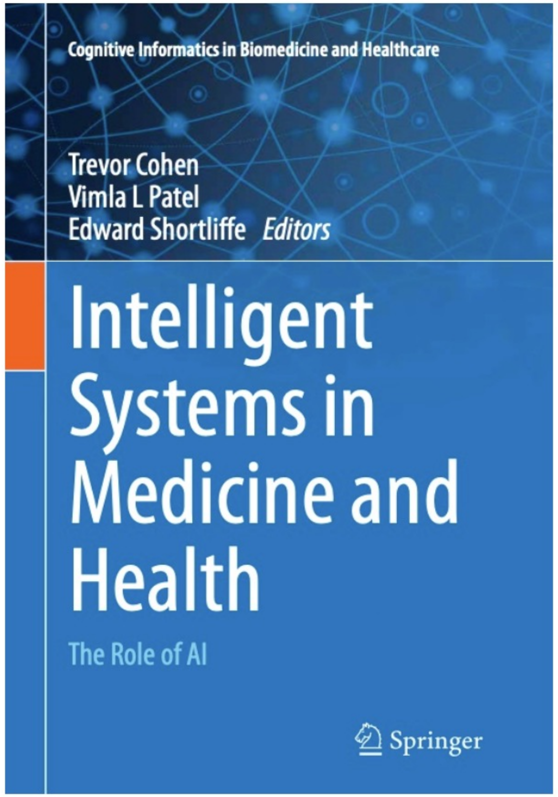

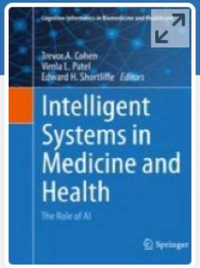

Intelligent Systems in Medicine and Health pp 343–371 | Cite as

Home > Intelligent Systems in Medicine and Health > Chapter

Interpreting Medical Images

Zongwei Zhou, Michael B. Gotway & Jianming Liang

Chapter | First Online: 10 November 2022

663 Accesses | 1 Citations

Part of the Cognitive Informatics in Biomedicine and Healthcare book series (CIBH)

Image-guided surgery	Quantitative analysis of disease progression	**Earlier detection of cancer**	Radiomics and predictive analytics	Neuroimaging

The FELIX Project

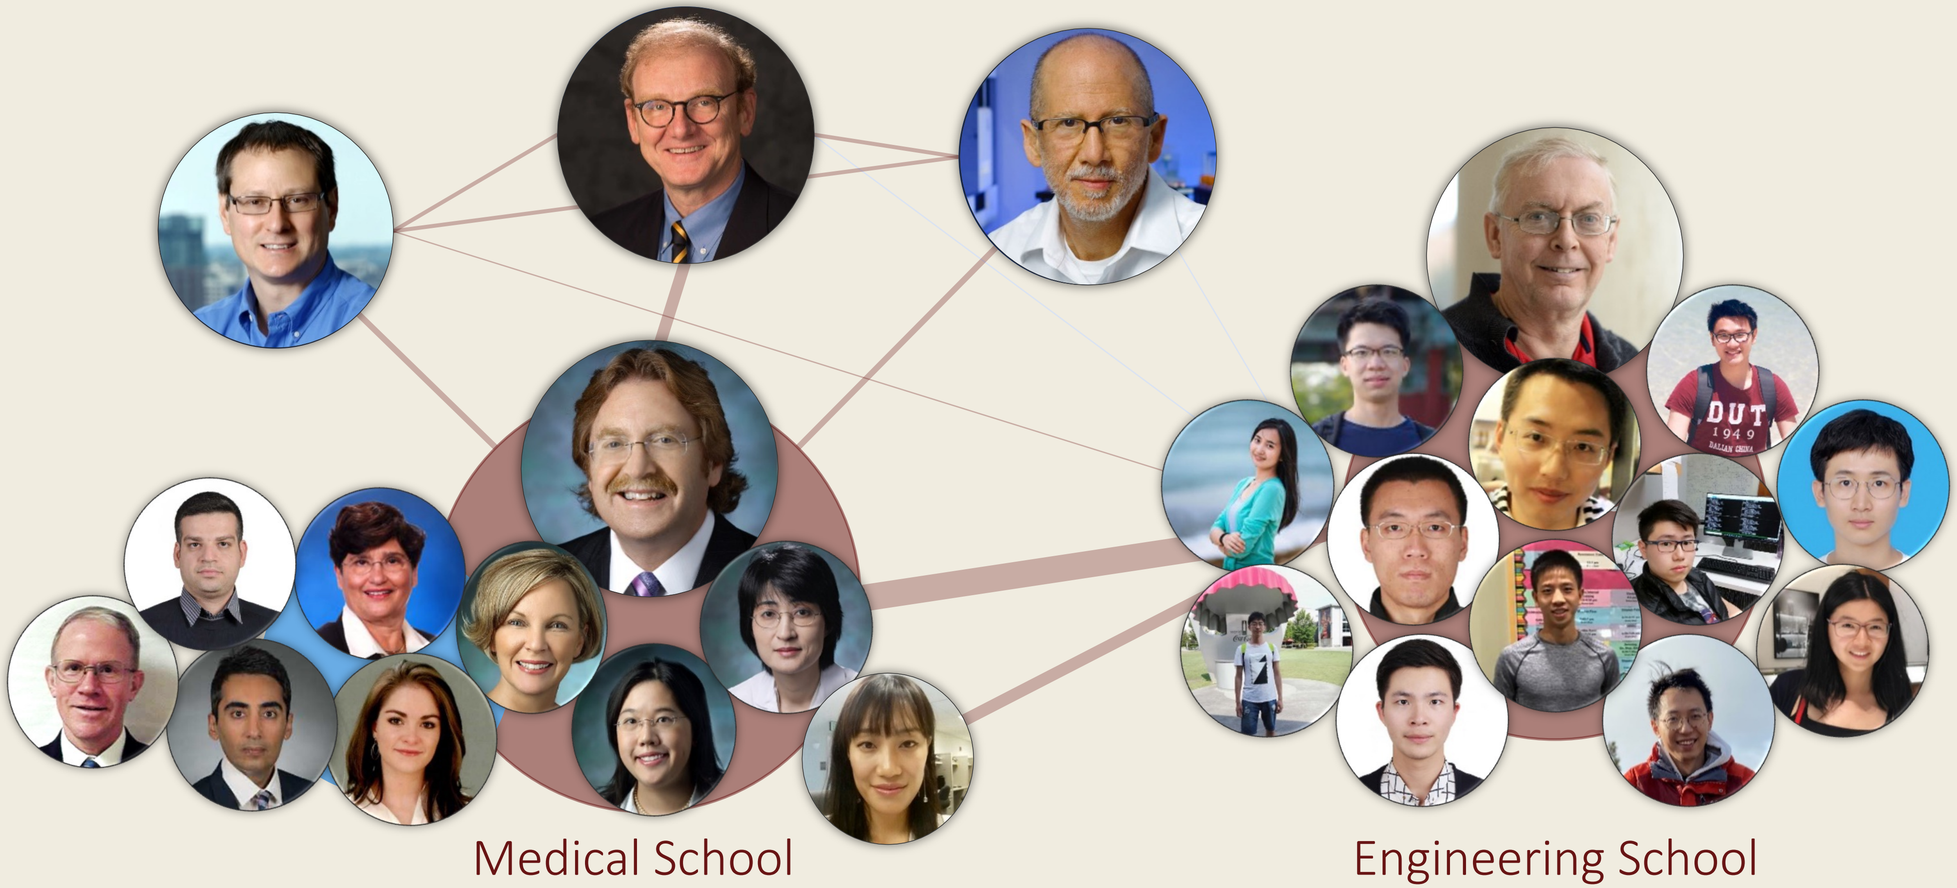

The FELIX Project

Goal: Earlier detection of pancreatic cancer

- 40,000,000 abdominal CT scans are performed each year in U.S.
- 1/3 of pancreatic cancer in these scans are missed by Radiologists.
- Pancreatic cancer is treatable if detected early.
- Deep Learning can see things in images that most humans miss.
    - 5,038 annotated CT scans at Johns Hopkins ☛ Sensitivity=97%, Specificity=99%
    - *This dataset took* ***15 years*** *to annotate for a human.*

1. Xia, Yingda, Qihang Yu, Linda Chu, Satomi Kawamoto, Seyoun Park, Fengze Liu, Jieneng Chen et al. "The felix project: Deep networks to detect pancreatic neoplasms." medRxiv (2022): 2022-09.

The FELIX-Civitas Project

Goal: Earlier detection of pancreatic cancer

New Goal: Earlier detection of a variety of cancers

- Body Maps: 3D Maps of Whole Body
  - Conceptually similar to Google Maps, but it focuses on human anatomy rather than the Earth's geography.
    - (1) Accurate segmentation of 104 anatomical structures.
    - (2) Cancer screening and localization across various structures.
    - (3) Language interaction between users and systems.

The FELIX-Civitas Project

Goal: Earlier detection of pancreatic cancer

New Goal: Earlier detection of a variety of cancers

- Body Maps: 3D Maps of Whole Body
  - Conceptually similar to Google Maps, but it focuses on human anatomy rather than the Earth's geography.
    - (1) Accurate segmentation of 104 anatomical structures.
    - (2) Cancer screening and localization across various structures.
    - (3) Language interaction between users and systems.
  - McGovern and Lustgarten (role: Team Investigator; status: **awarded**)
  - NIH K99/R00 (role: PI; status: under review)
  - NIH R01 and ACS Grant (role: Team Investigator; status: under review)

Body Maps: 3D Maps of Whole Body

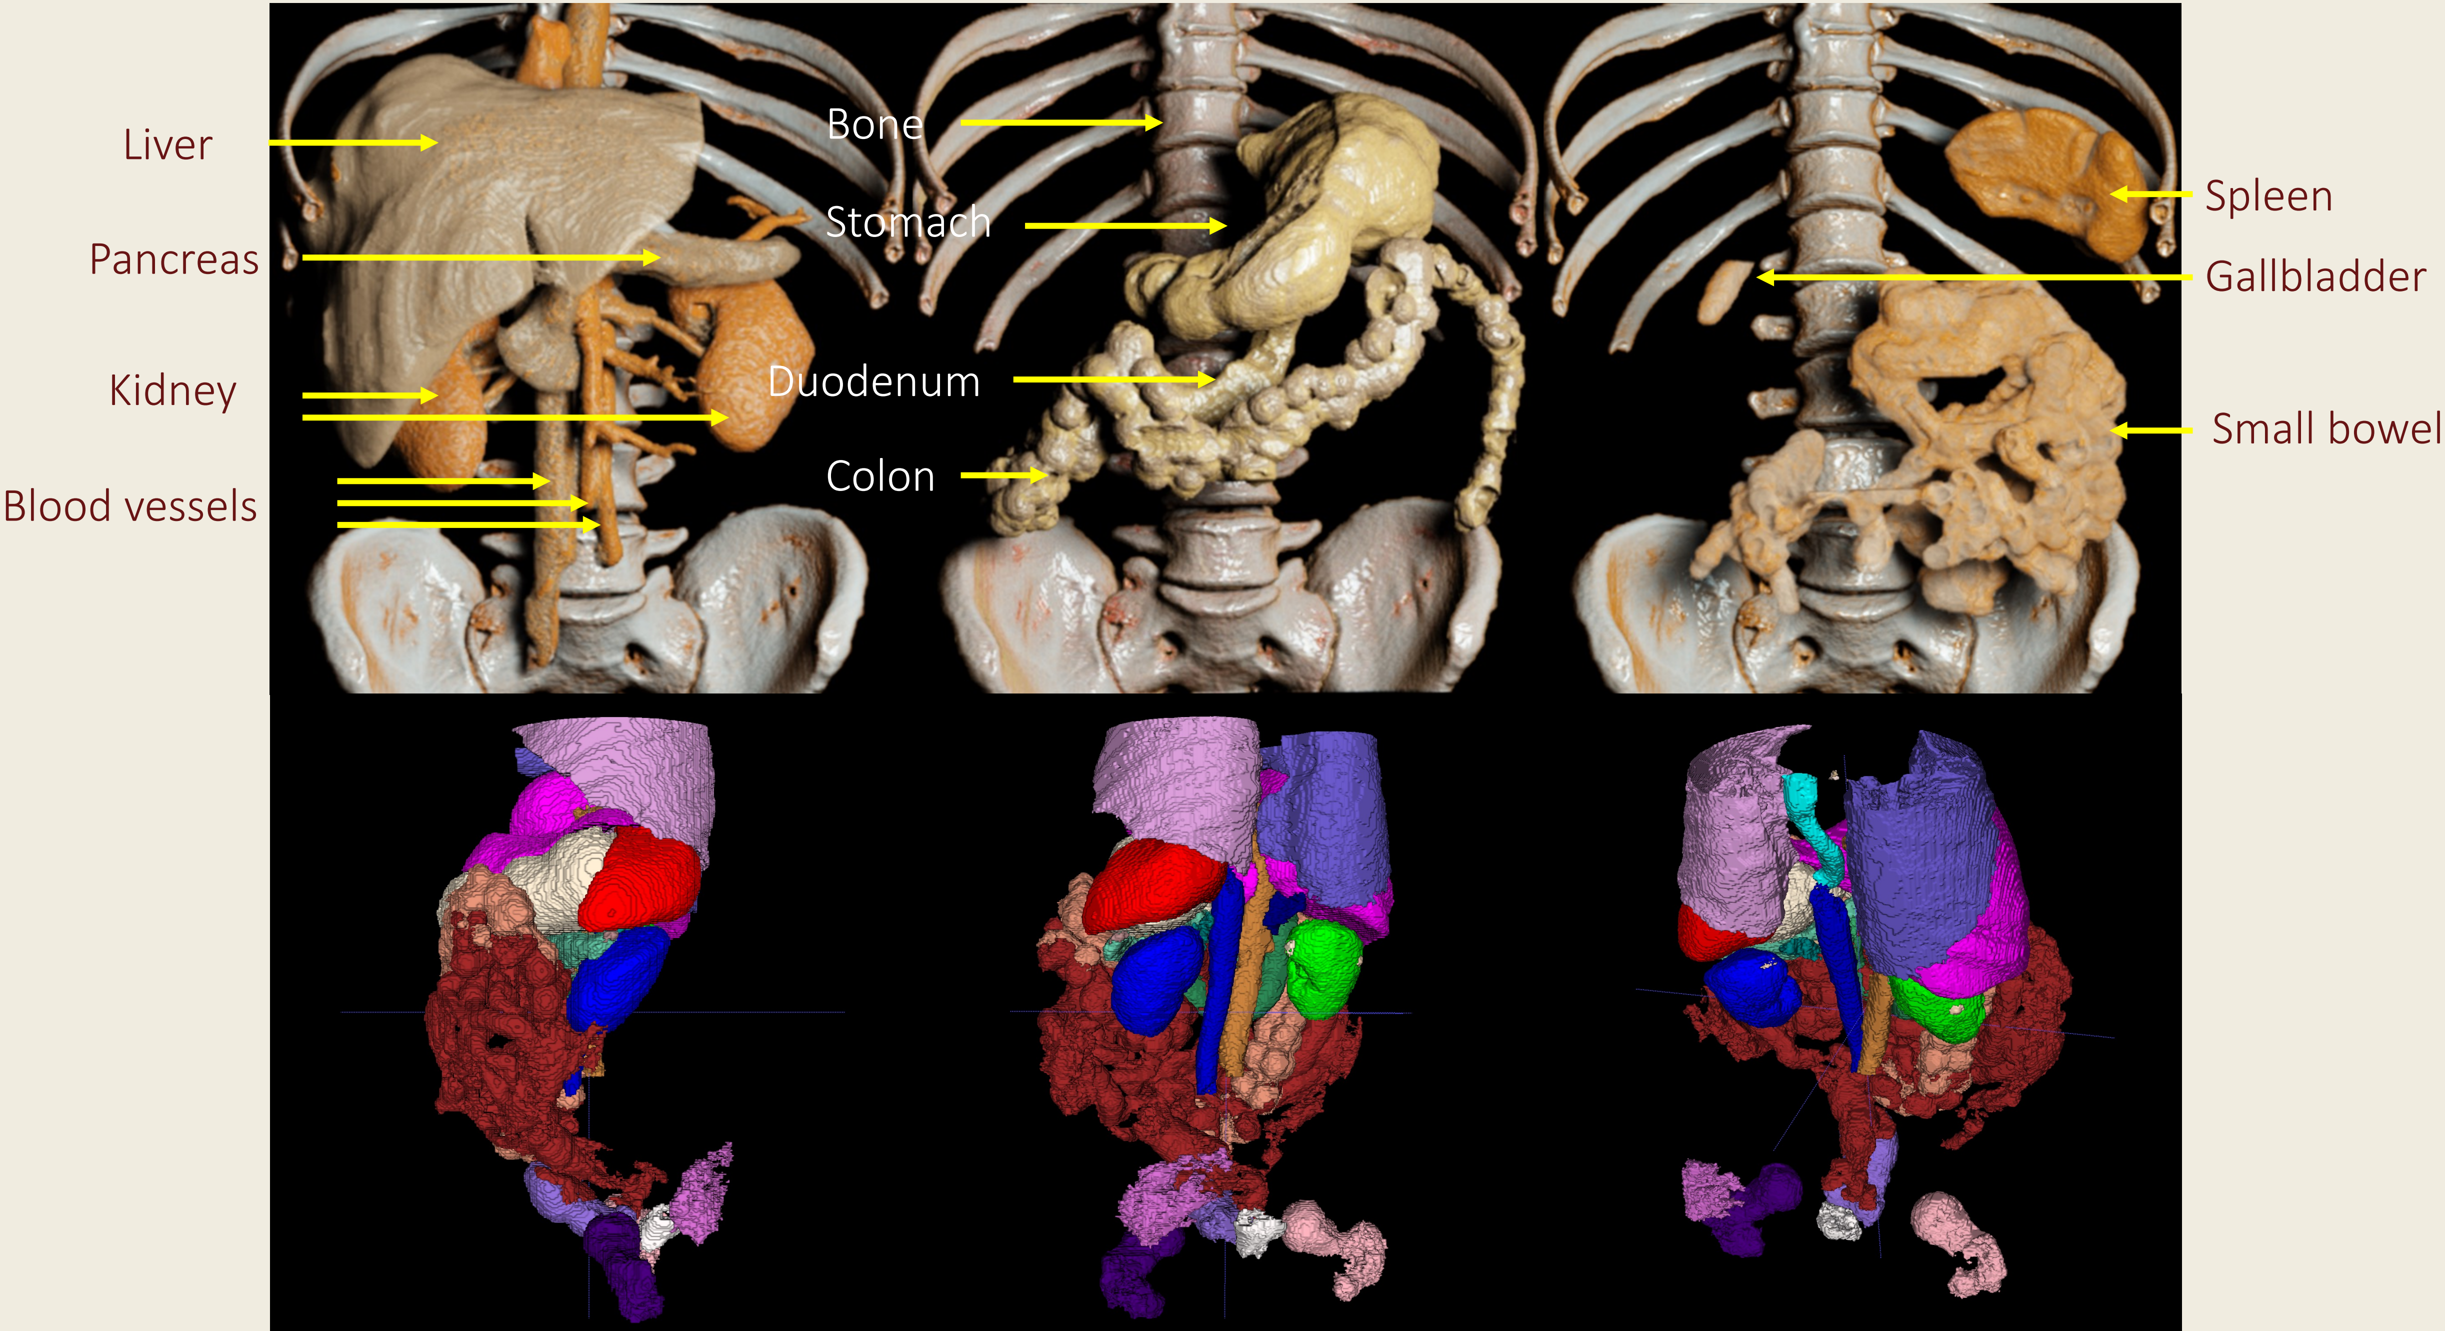

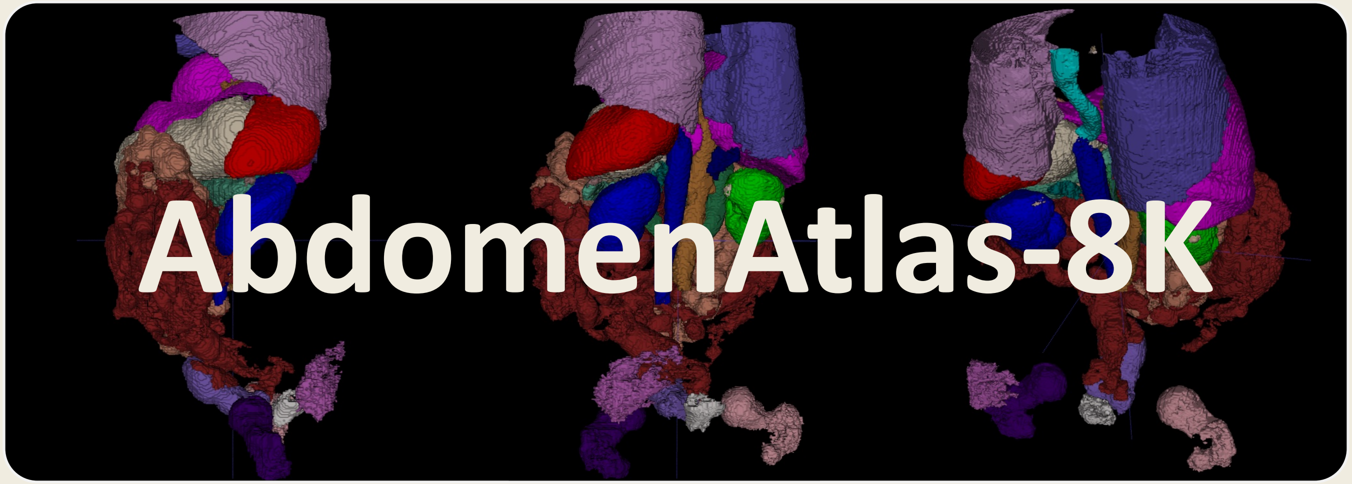
AbdomenAtlas-8K

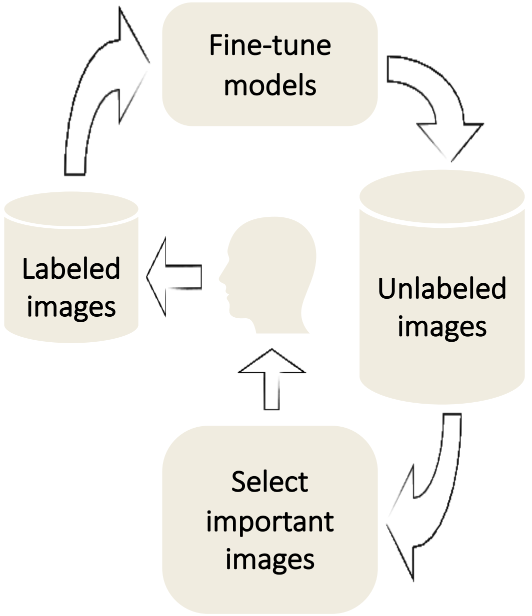

Interactive segmentation

Annotated
25
organs

Annotated
7
cancers

Integrated
15
public datasets

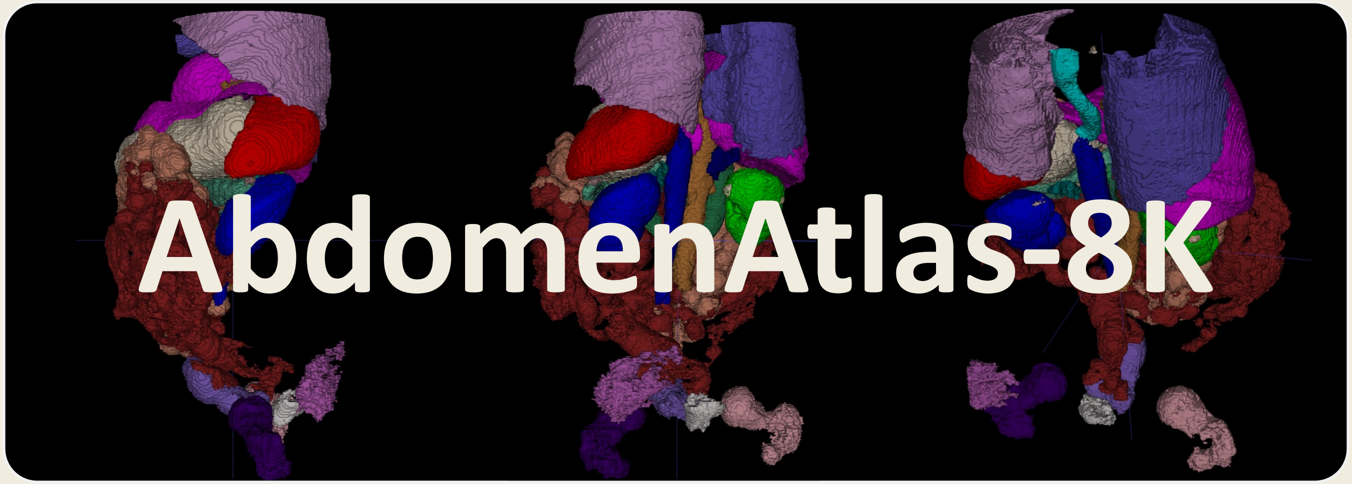

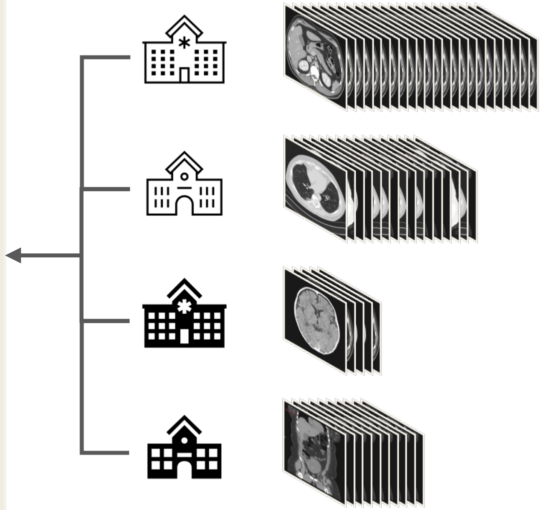

Collected from
27 hospitals
worldwide

Up to
533x faster
than previous strategies

MONAI

Annotated
3.2M
images

Annotated
8,448
CT volumes

Created in
3 Weeks
by 1 annotator

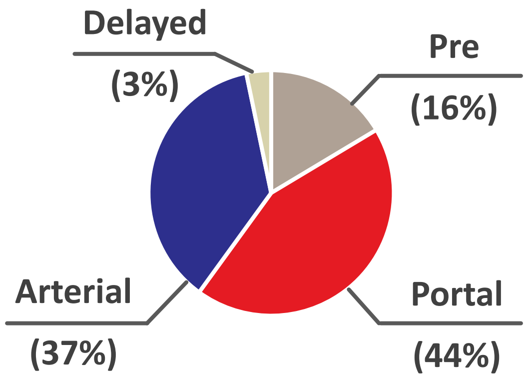

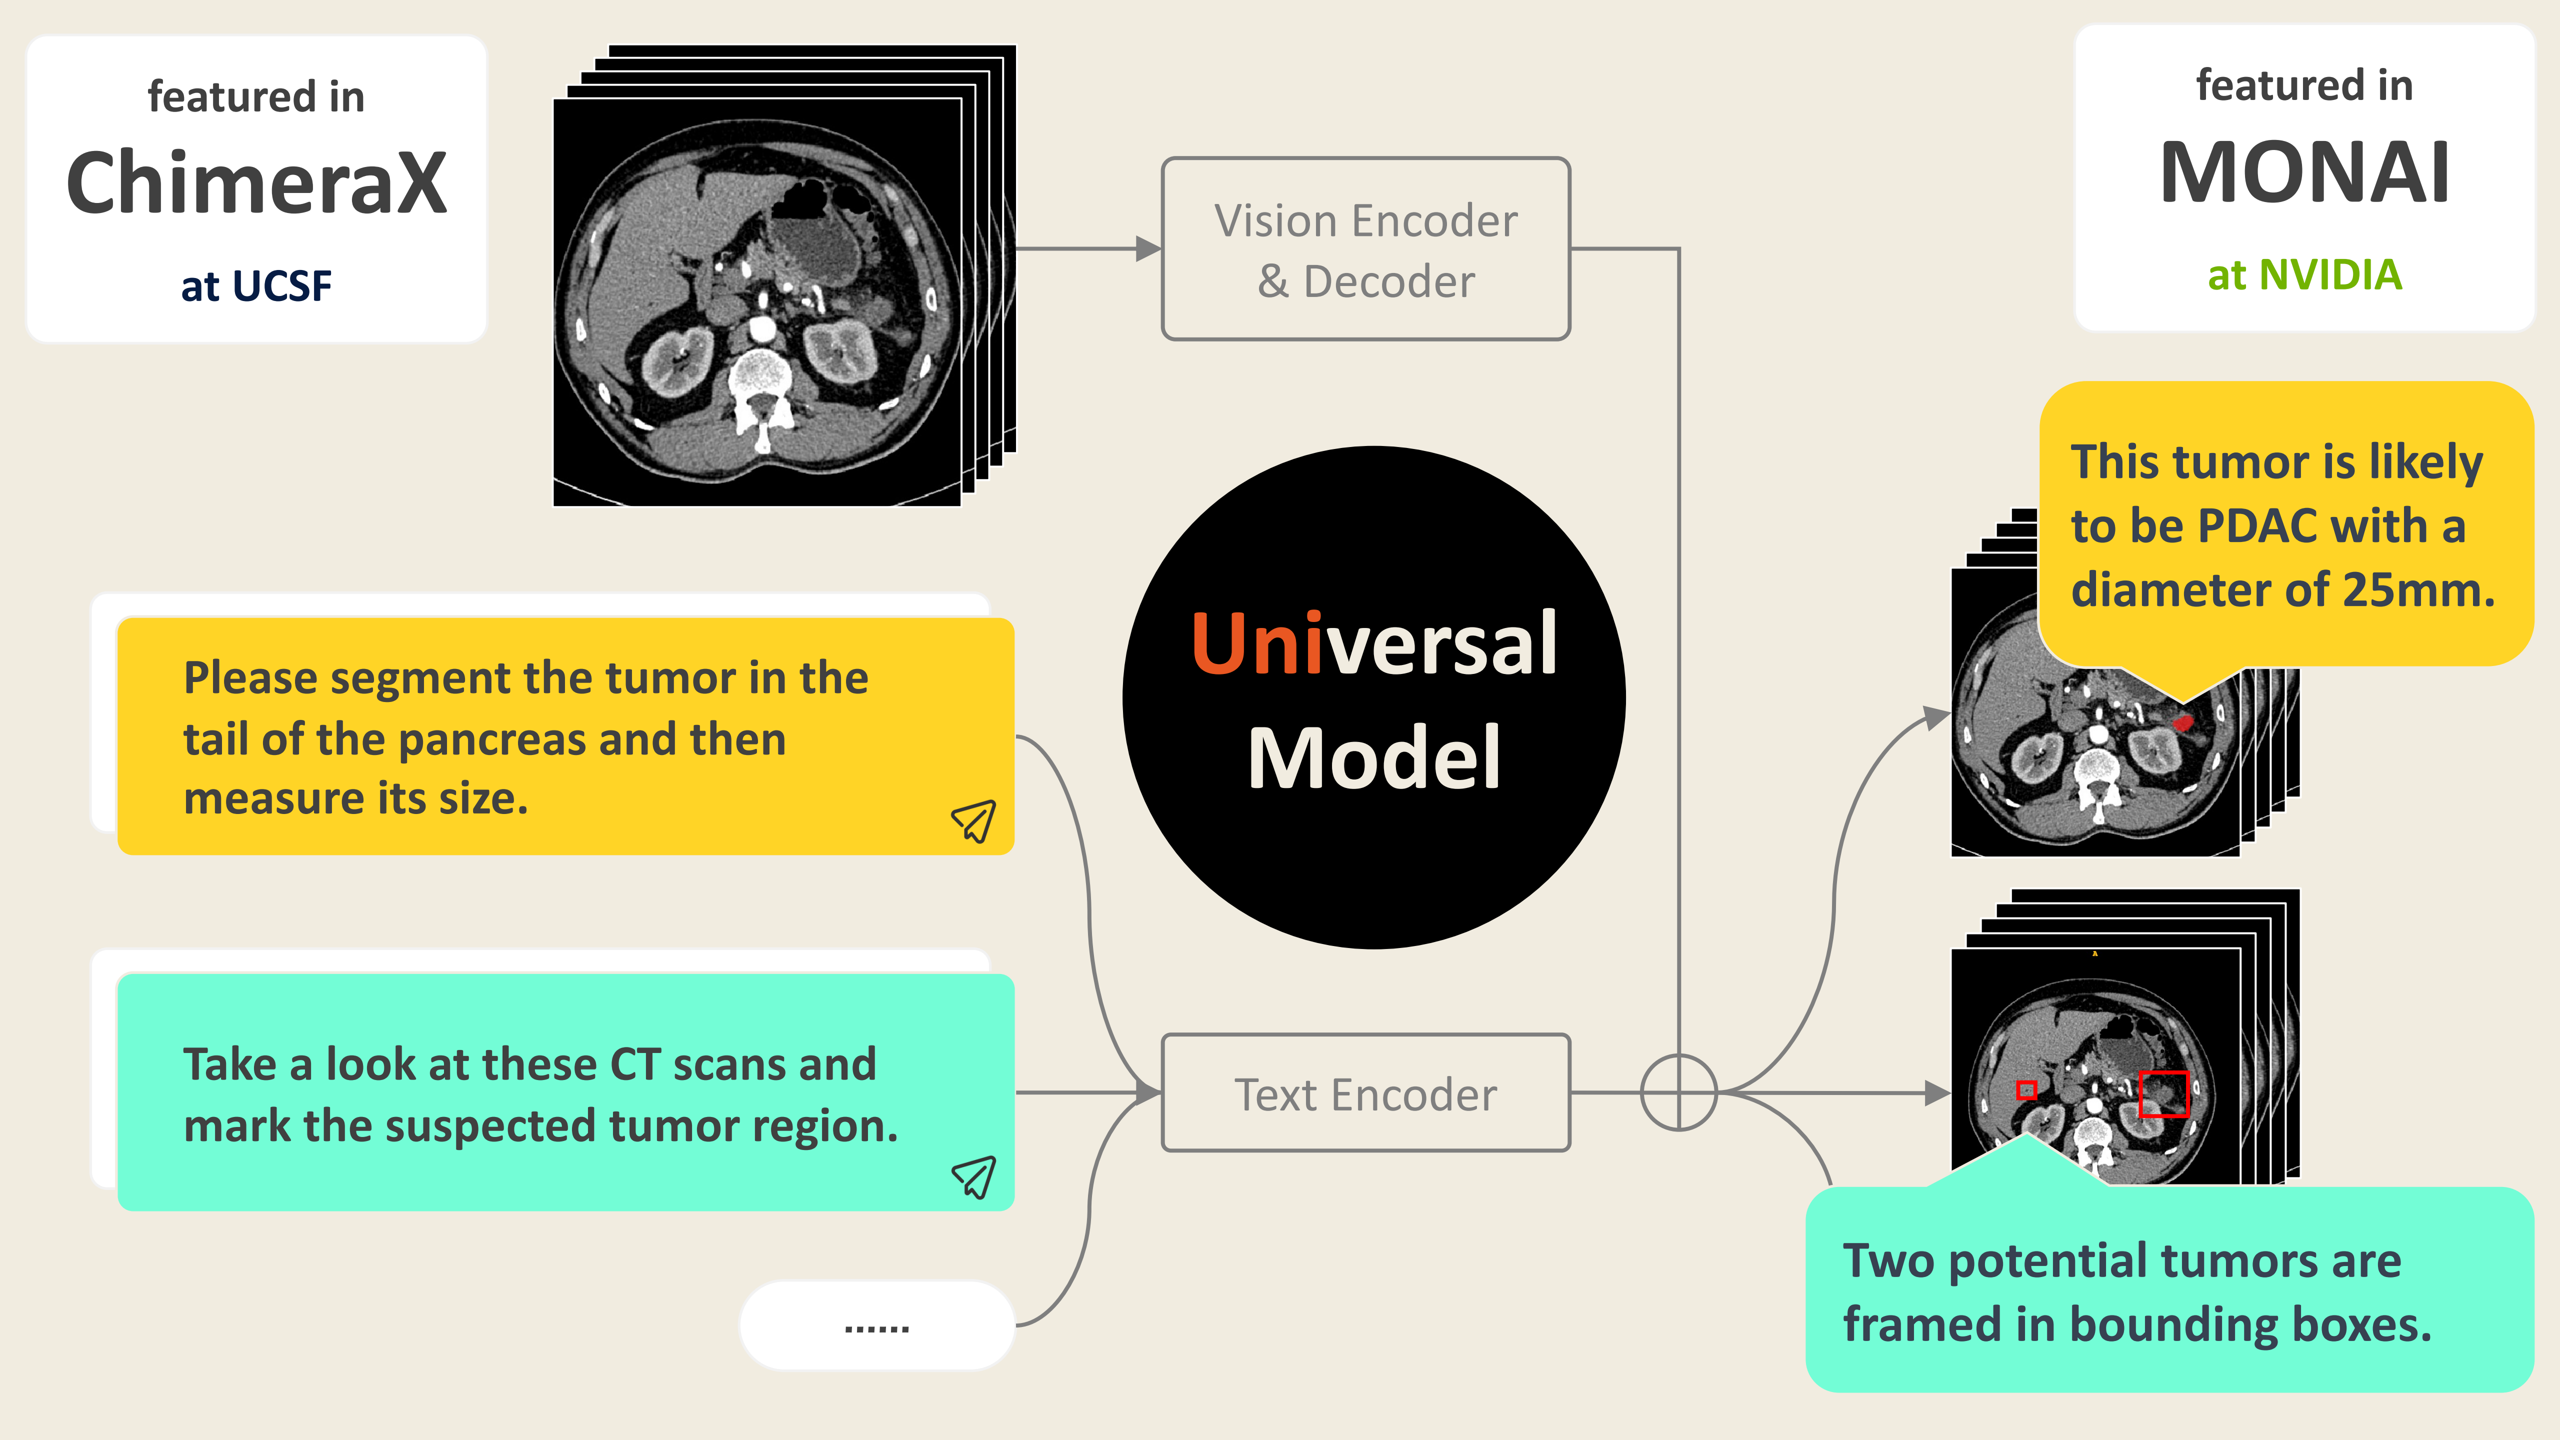
featured in
ChimeraX
at UCSF
featured in
MONAI
at NVIDIA
Vision Encoder & Decoder
Universal Model
Please segment the tumor in the tail of the pancreas and then measure its size.
Take a look at these CT scans and mark the suspected tumor region.
......
Text Encoder
This tumor is likely to be PDAC with a diameter of 25mm.
Two potential tumors are framed in bounding boxes.

AbdomenAtlas-8K

8,848 annotated CT volumes

Universal Model

25 organs and 7 cancers

Medical Segmentation Decathlon

Info | Leaderboard | Statistics | Join

#	User (Team)	Created	Mean Position
1st	zongwei.zhou (universal_model)	13 Feb. 2023	5.6
2nd	Swin_UNETR	13 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

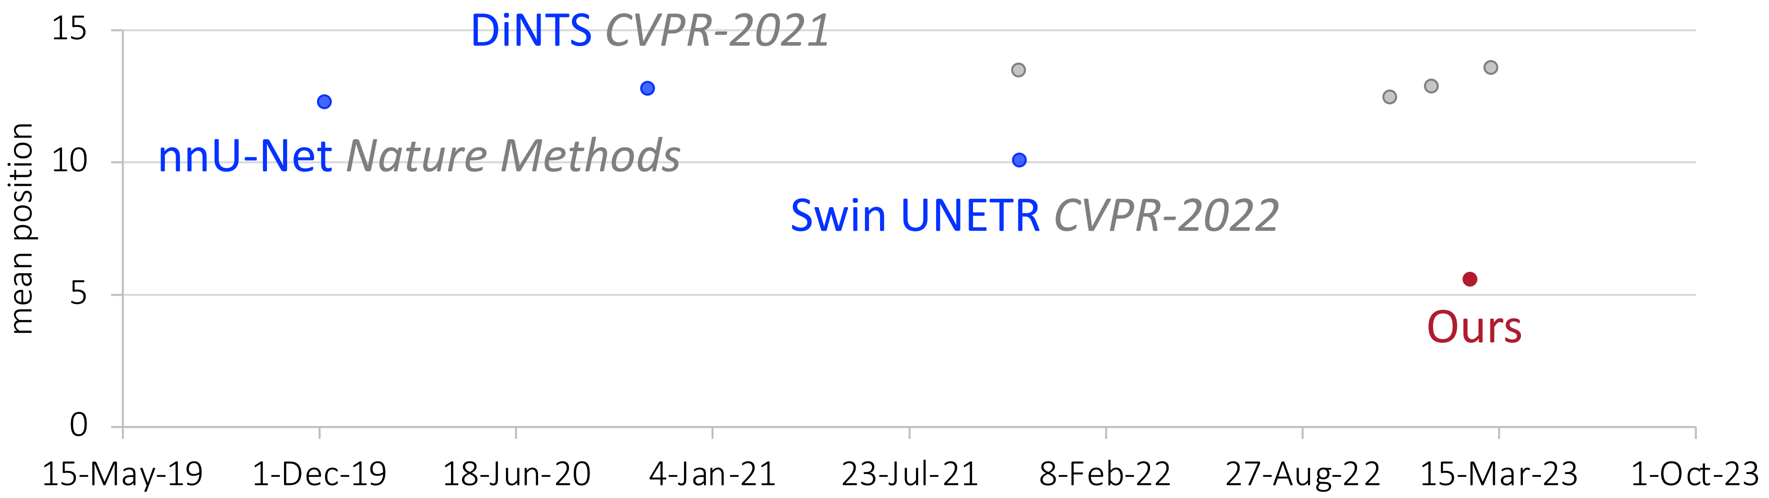

Earlier detection of cancers

Earlier detection of cancer

Developing 3D Maps of Whole Body

Significant

Applications

Impactful

Methodologies

Reducing Annotation Efforts for Radiologists

Annotation-intensive deep learning

Annotation-efficient deep learning

Annotation-free deep learning

New Chapter (2020s)

Challenges and Questions

Significant

Applications

Methodologies

Impactful

Scaling datasets

Multiple modalities
Diverse institutes
IRB approval

Tumor Synthesis

Annotation-free deep learning

Scaling annotations

Efficient annotation
Human in the loop
Novel disease

AbdomenAtlas-8K

8,848 annotated CT volumes

Scaling algorithms

Vision-language
Lifelong learning
Reader study

Universal Model

25 organs and 7 cancers

Thank You!

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

Statistics
7,118 citations
Top 2% of Scientists in 2022
24 first/corresponding authored papers