Cost-Effective Deep Learning in Medical Image Analysis

Zongwei Zhou

College of Health Solutions, Arizona State University

P: 1-(480)738-2575 | E: zongweiz@asu.edu

Introduction
Significance
Aim #1
Aim #2
Aim #3
Summary

REVIEW

Deep learning

Yann LeCun[1,2], Yoshua Bengio[3] & Geoffrey Hinton[4,5]

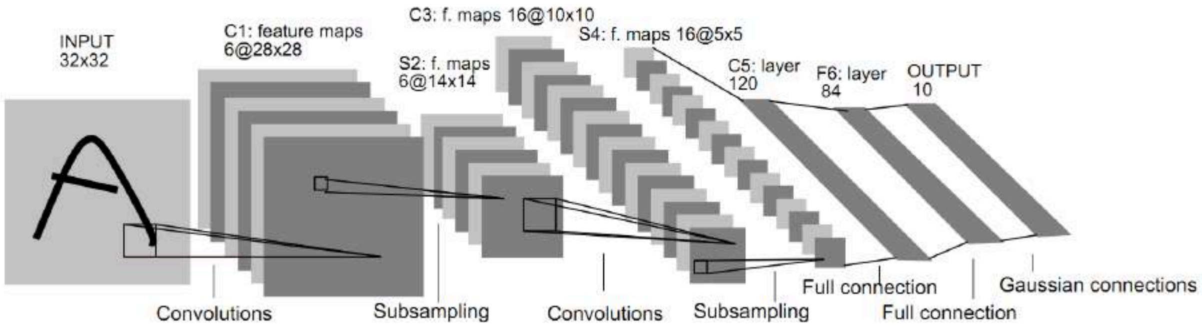

Turing Award And $1 Million Given To 3 AI Pioneers

Nicole Martin Former Contributor
AI & Big Data
I write about digital marketing, data and privacy concerns.

Advantages: Generalizable

- Lung cancer
- Skin cancer
- Diabetic retinopathy

Barrier to medical imaging:

- Large annotation cost

Introduction
Significance
Aim #1
Aim #2
Aim #3
Summary

LETTERS
https://doi.org/10.1038/s41591-019-0447-x
nature medicine
Corrected: Author Correction

End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography

Diego Ardila[1,5], Atilla P. Kiraly[1,5], Sujeeth Bharadwaj[1,5], Bokyung Choi[1,5], Joshua J. Reicher[2], Lily Peng[1], Daniel Tse[1*], Mozziyar Etemadi[3], Wenxing Ye[1], Greg Corrado[1], David P. Naidich[4] and Shravya Shetty[1]

42,290

CT images

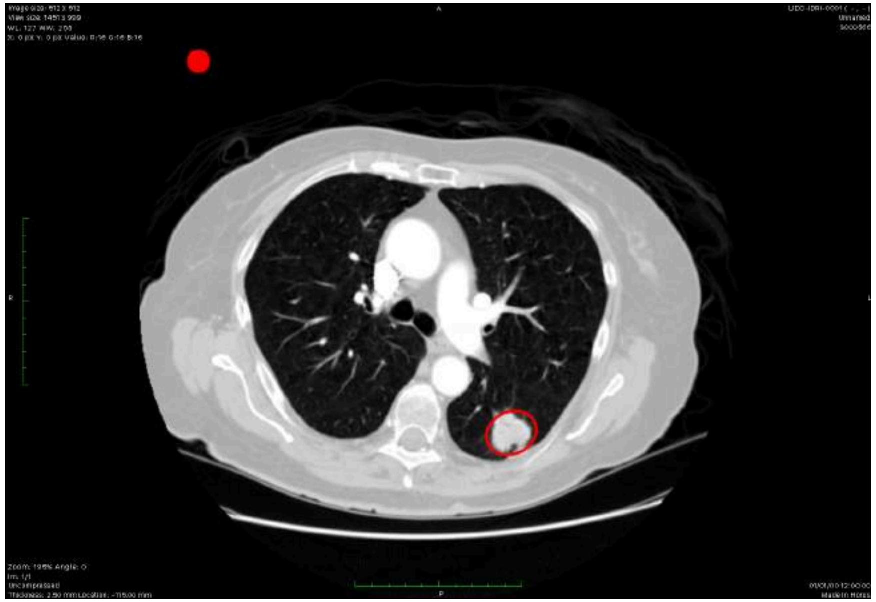

Introduction

Significance

Aim #1

Aim #2

Aim #3

Summary

LETTER

doi:10.1038/nature21056

Dermatologist-level classification of skin cancer with deep neural networks

Andre Esteva[1]*, Brett Kuprel[1]*, Roberto A. Novoa[2,3], Justin Ko[2], Susan M. Swetter[2,4], Helen M. Blau[5] & Sebastian Thrun[6]

129,450

clinical images

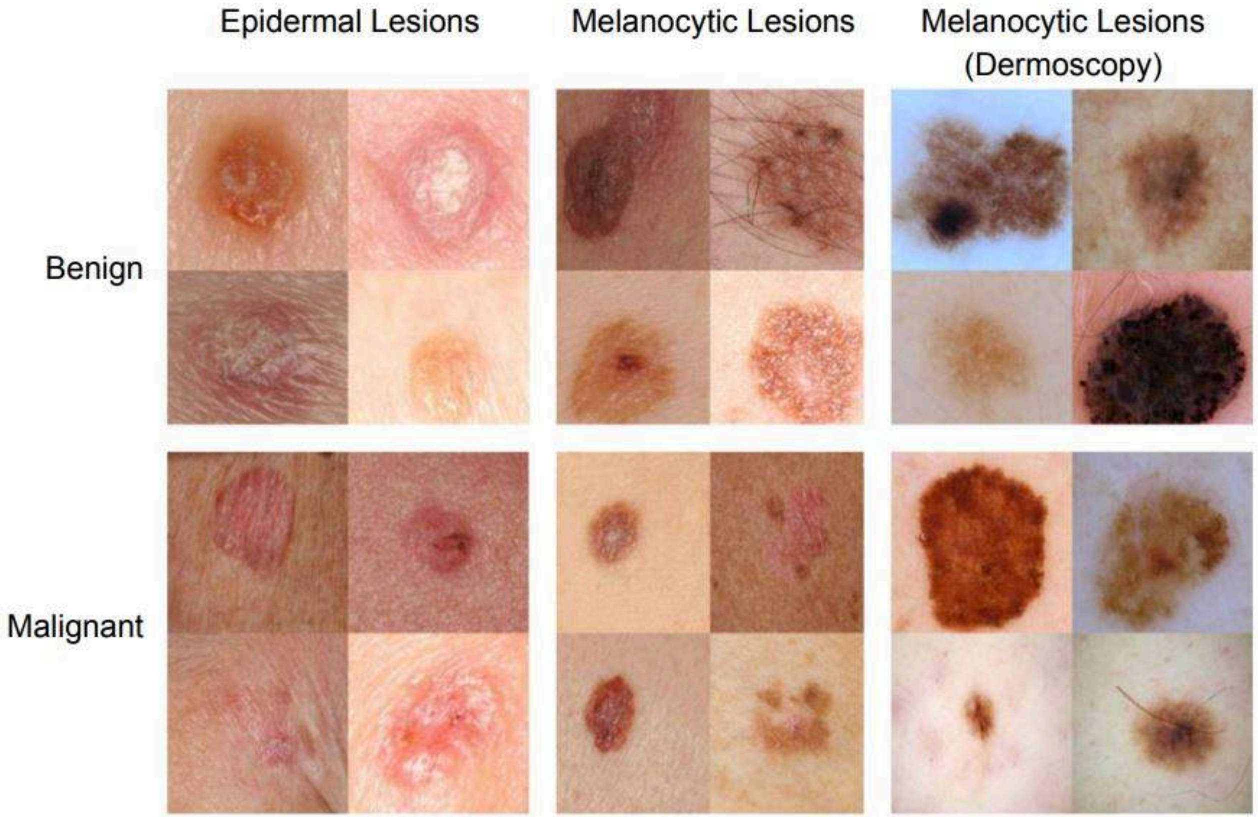

Introduction

Significance

Aim #1

Aim #2

Aim #3

Summary

Research

JAMA | Original Investigation | INNOVATIONS IN HEALTH CARE DELIVERY

Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs

Varun Gulshan, PhD; Lily Peng, MD, PhD; Marc Coram, PhD; Martin C. Stumpe, PhD; Derek Wu, BS; Arunachalam Narayanaswamy, PhD; Subhashini Venugopalan, MS; Kasumi Widner, MS; Tom Madams, MEng; Jorge Cuadros, OD, PhD; Ramasamy Kim, OD, DNB; Rajiv Raman, MS, DNB; Philip C. Nelson, BS; Jessica L. Mega, MD, MPH; Dale R. Webster, PhD

128,175

retinal images

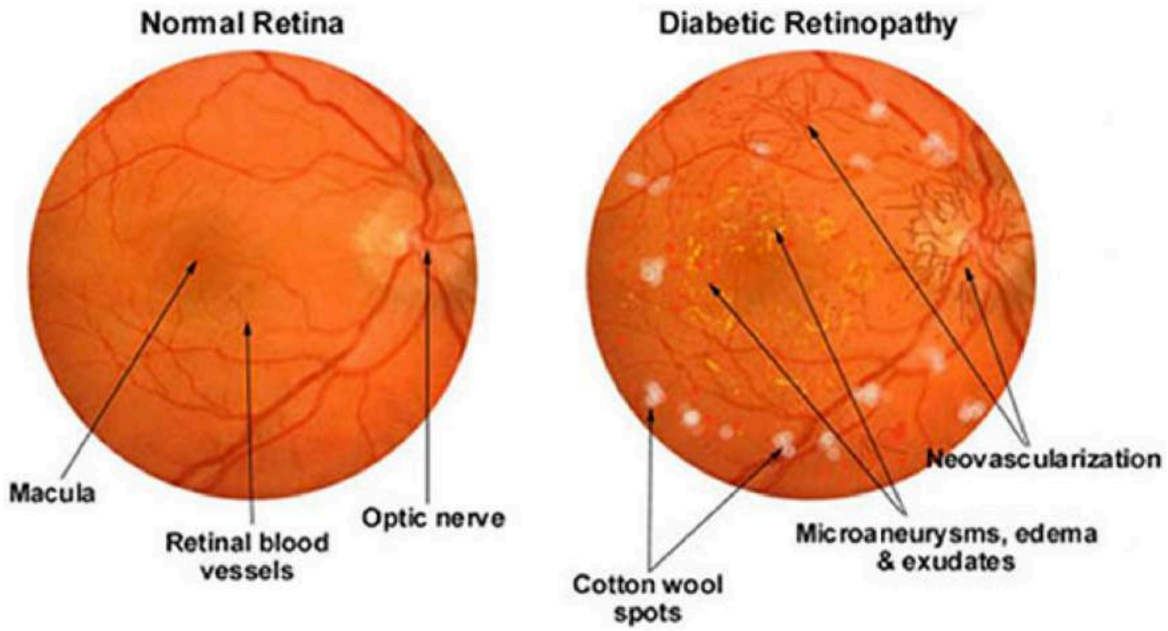

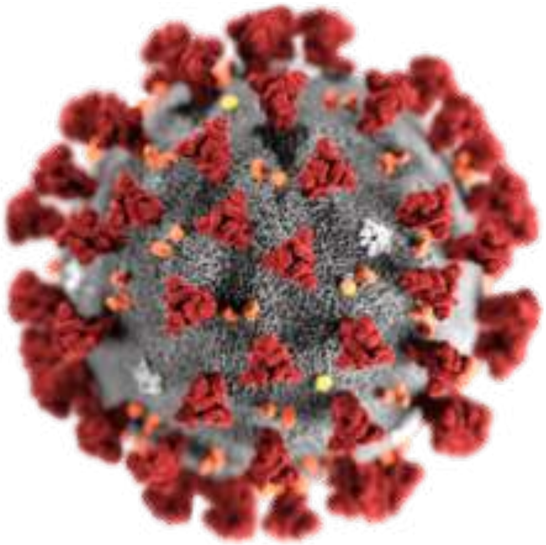

COVID-19?

How to develop an effective deep learning algorithms for those diseases that have no such labeled big data?

Quickly develop an effective computer-aided diagnosis system is important

- A flood of patients are pending during an outbreak
- Doctors do not have time to annotate every case
- Not many doctors have expertise for novel diseases

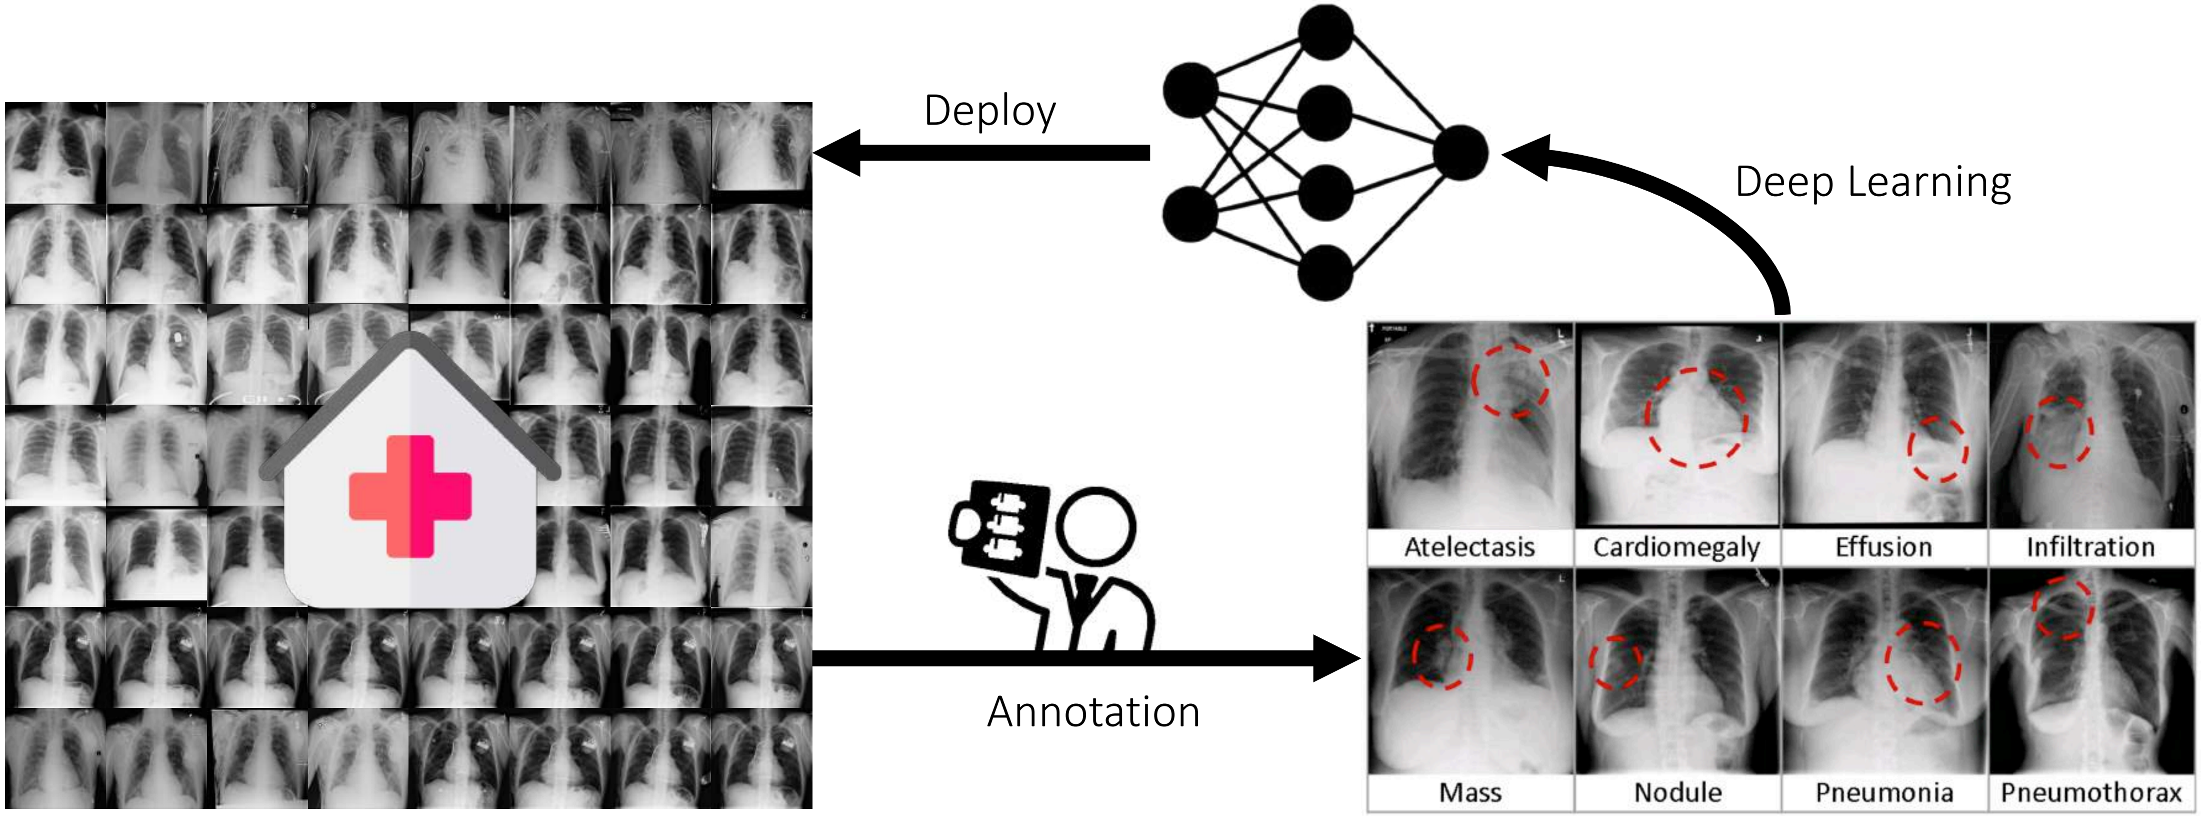

$1,000,000 annotation cost ☹

$1,000 annotation budget ☺

Research goal: Exploit novel methods to minimize the manual labeling efforts for a rapid, precise computer-aided diagnosis system

Aim #1: Acquire necessary annotation efficiently from human experts

Aim #2: Utilize existing annotation effectively from advanced architecture

Aim #3: Extract generic knowledge directly from unannotated images

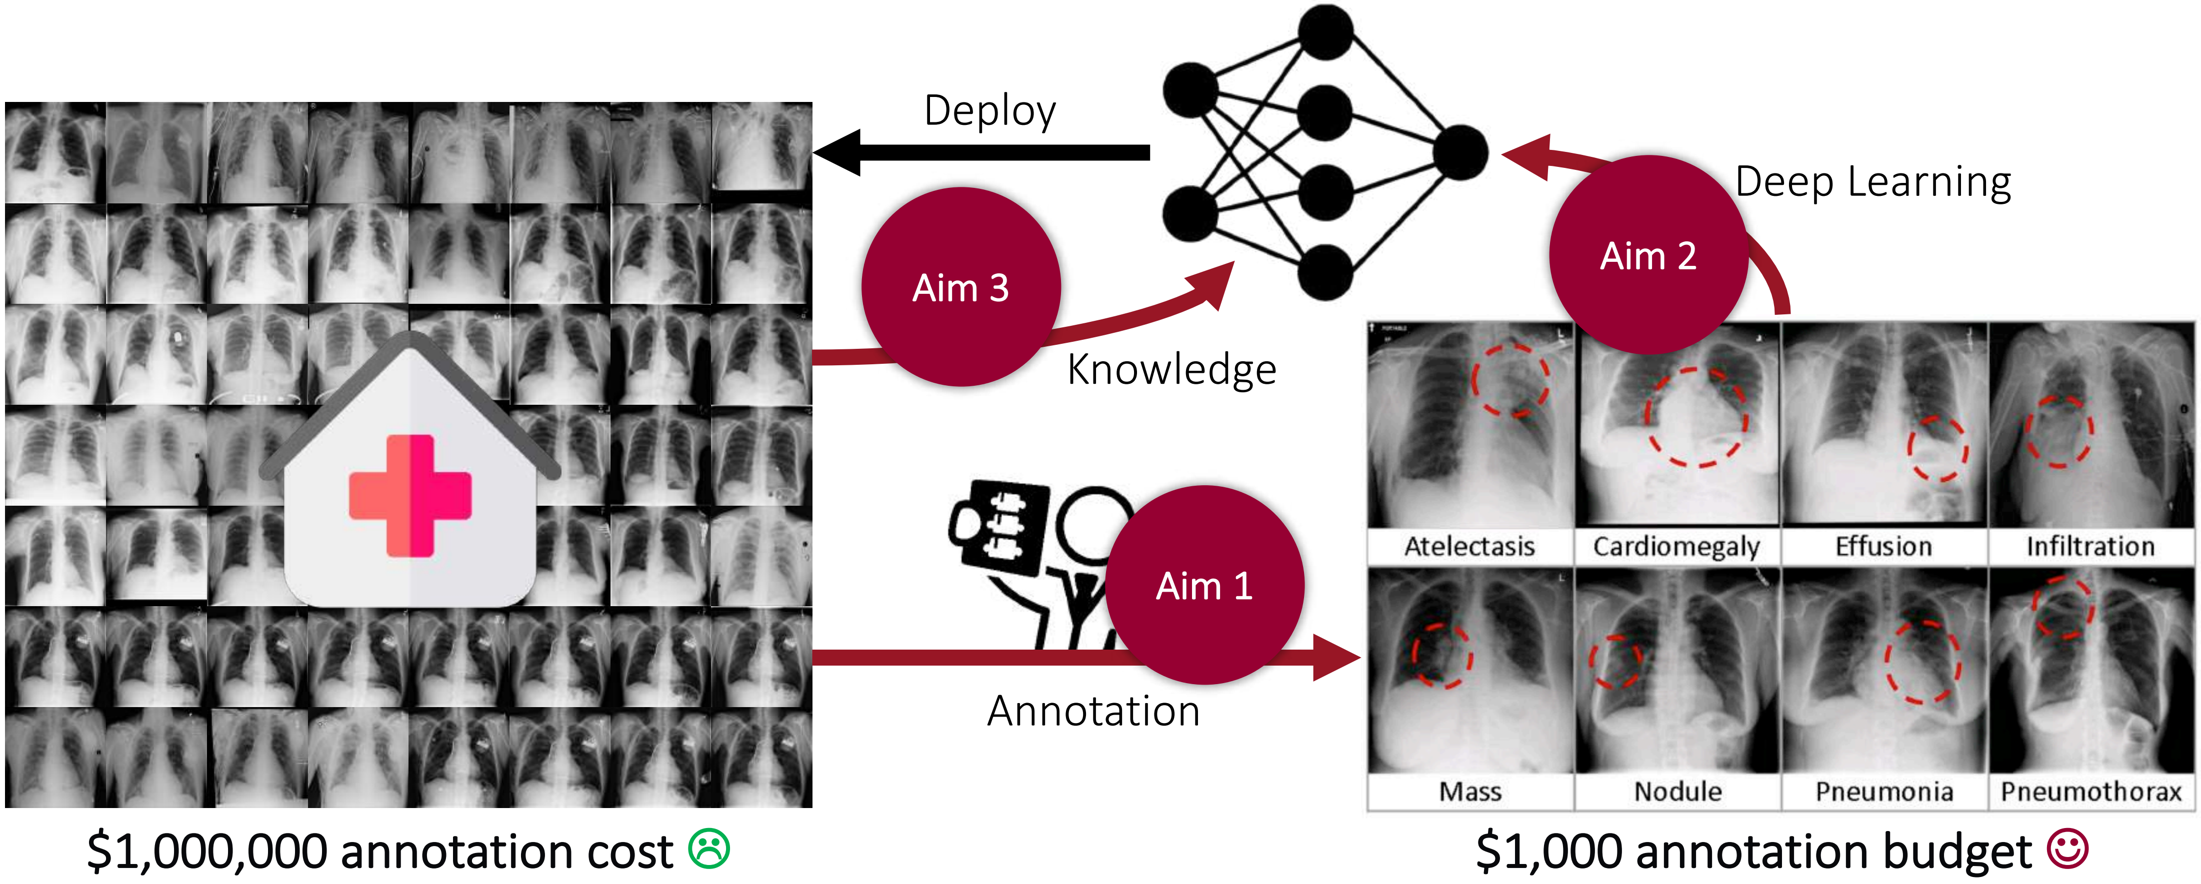

Research goal: Exploit novel methods to minimize the manual labeling efforts for a rapid, precise computer-aided diagnosis system

Aim #1: Acquire necessary annotation efficiently from human experts

Aim #2: Utilize existing annotation effectively from advanced architecture

Aim #3: Extract generic knowledge directly from unannotated images

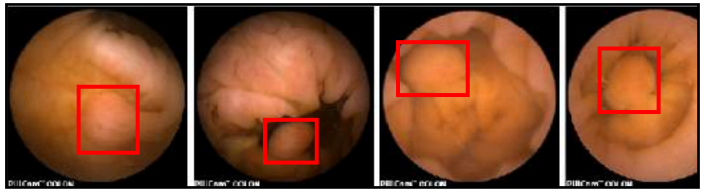
Polyp detection

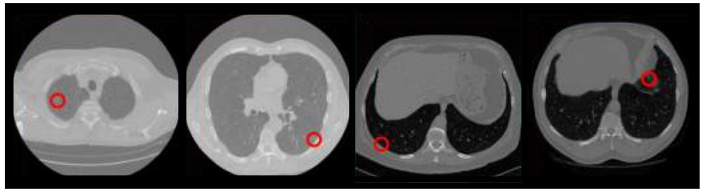
Lung nodule detection

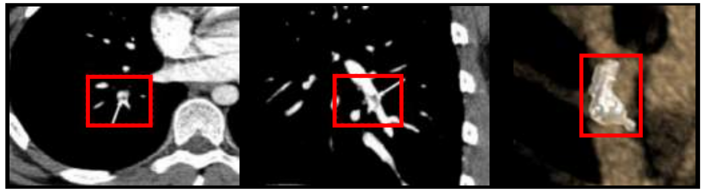
Pulmonary embolism detection

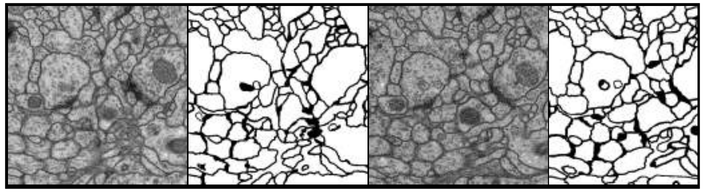
Neuronal structure segmentation

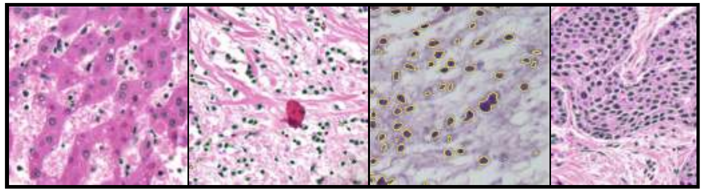
Cell/nuclei segmentation

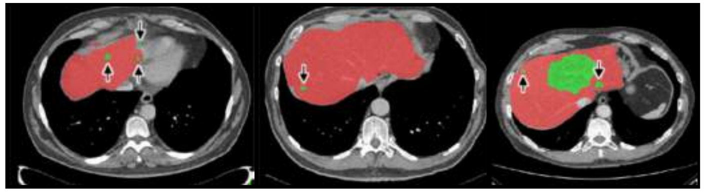
Liver/lesion segmentation

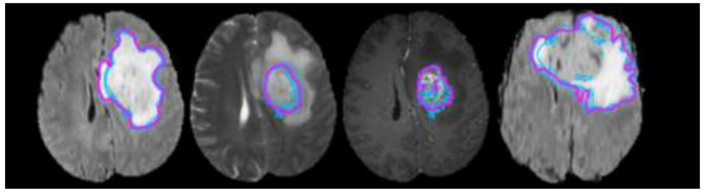
Brain/tumor segmentation


Kidney/lesion segmentation

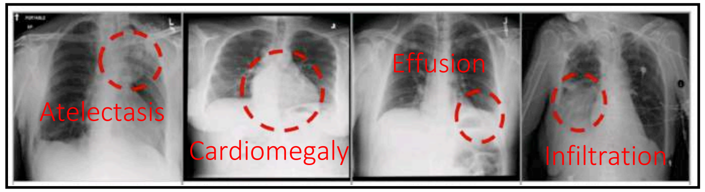

Pulmonary diseases classification

Aim #1: Acquire necessary annotation efficiently from human experts

Problem: Find the most important 1,000 images from 1,000,000 images

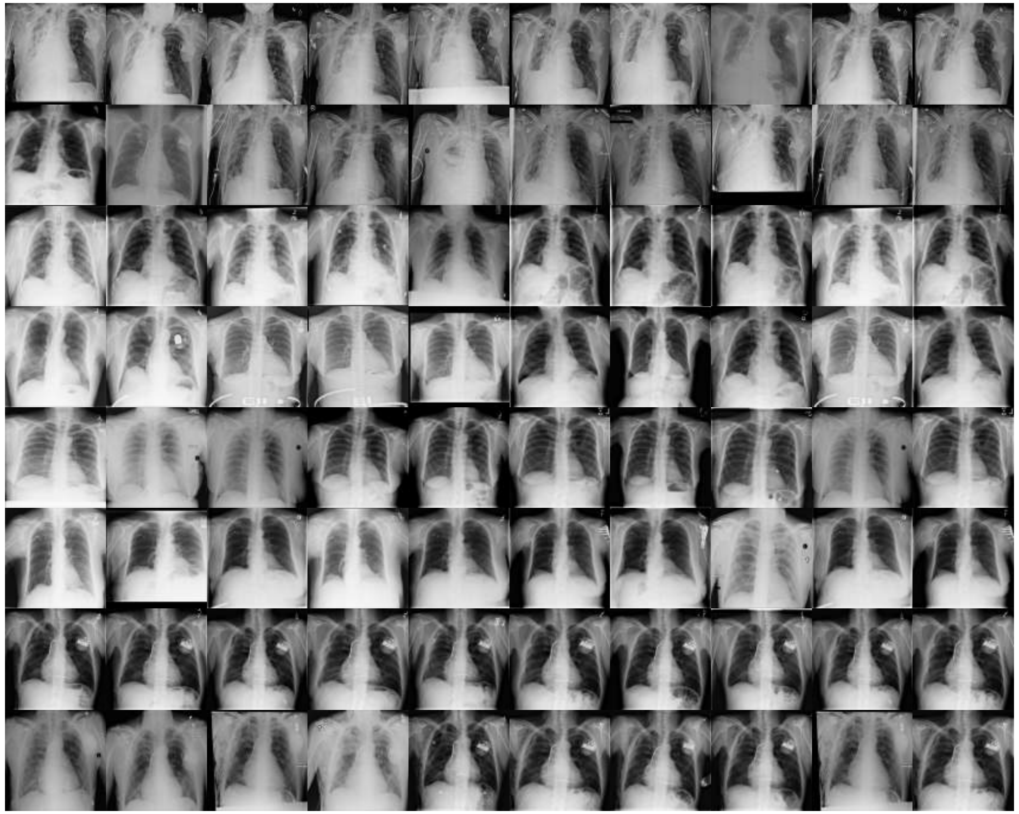

$ 1,000,000 annotation cost ☹

$1 per subject

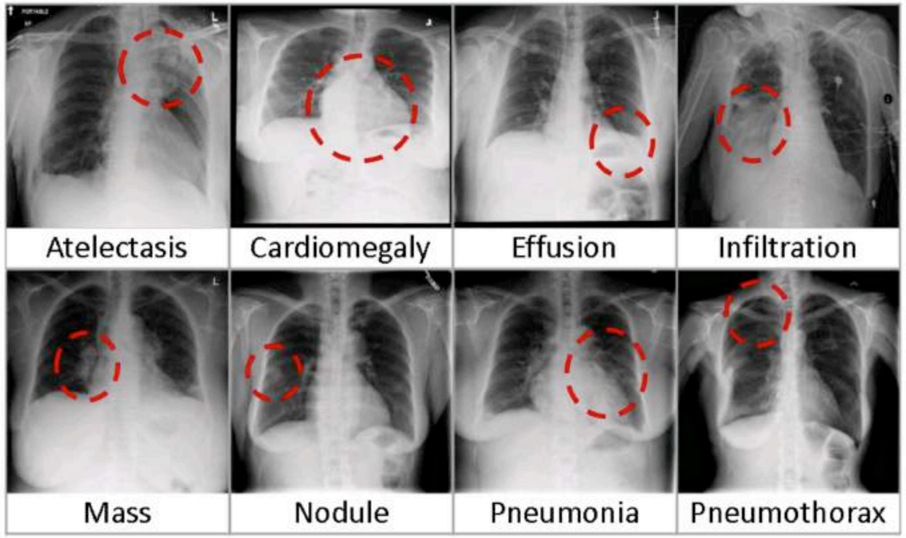

$ 1,000 annotation budget ☺

Introduction

Significance

Aim #1

Aim #2

Aim #3

Summary

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

Pre-trained models

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

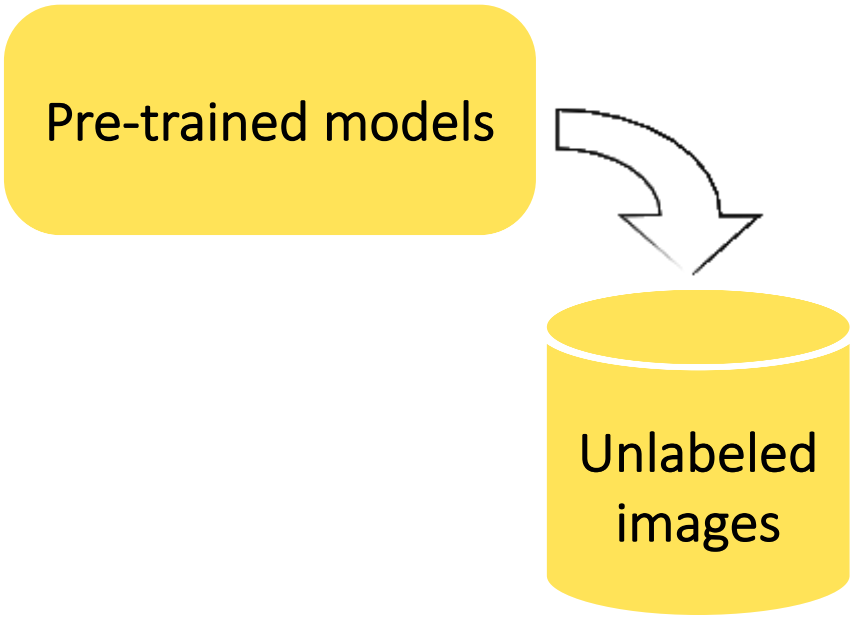

Introduction

Significance

Aim #1

Aim #2

Aim #3

Summary

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

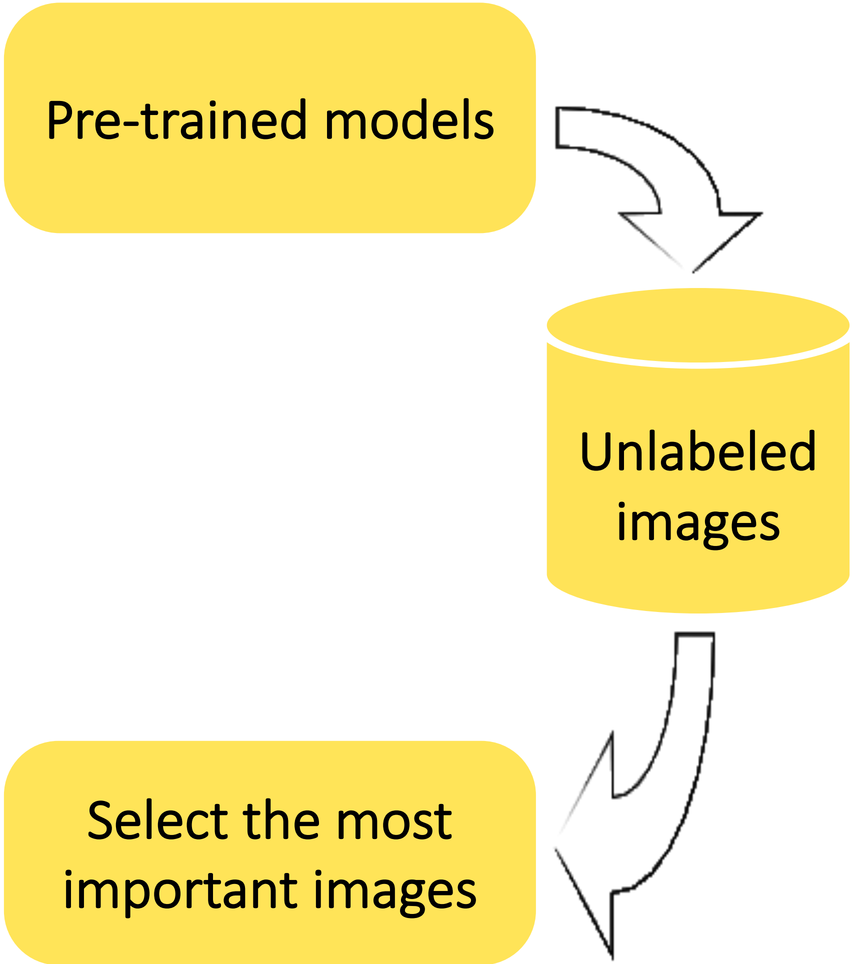

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

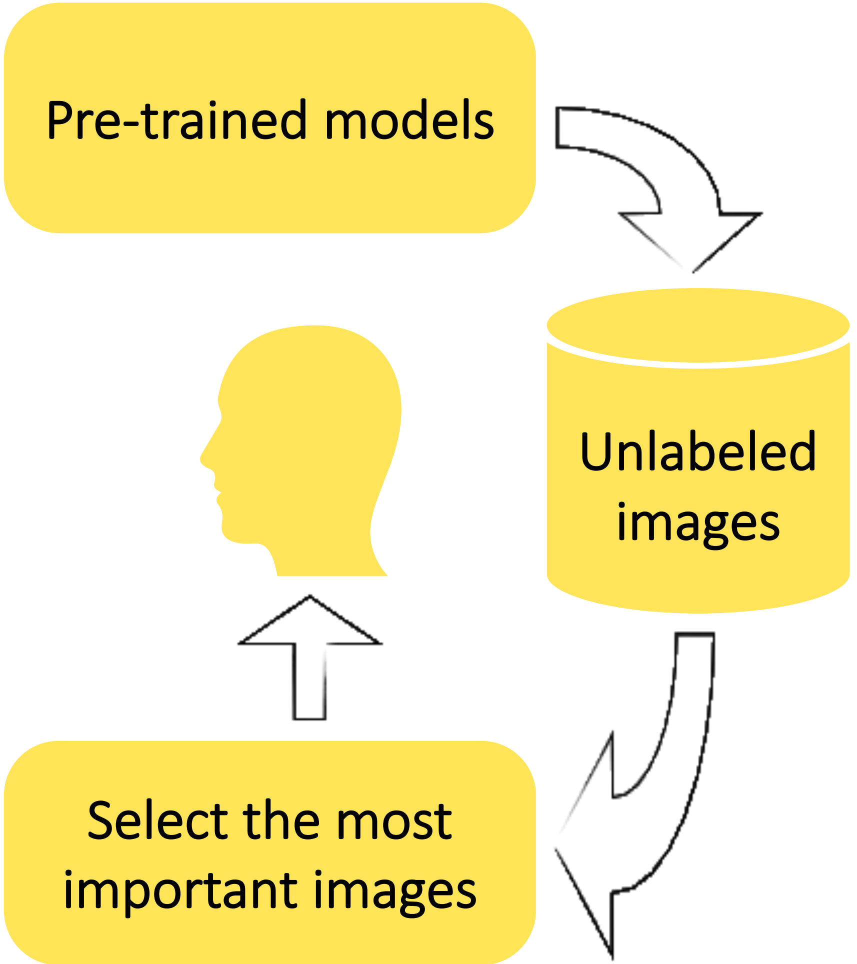

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

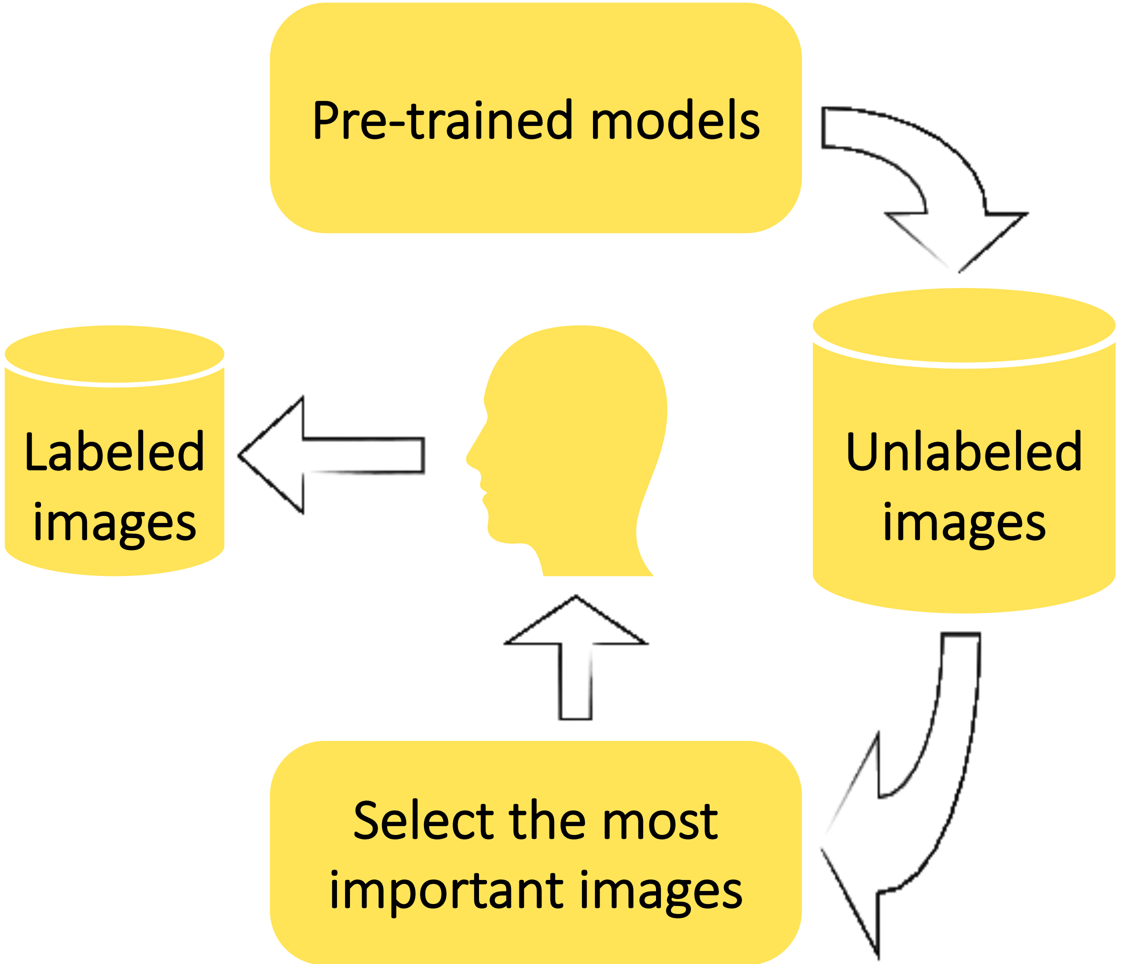

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

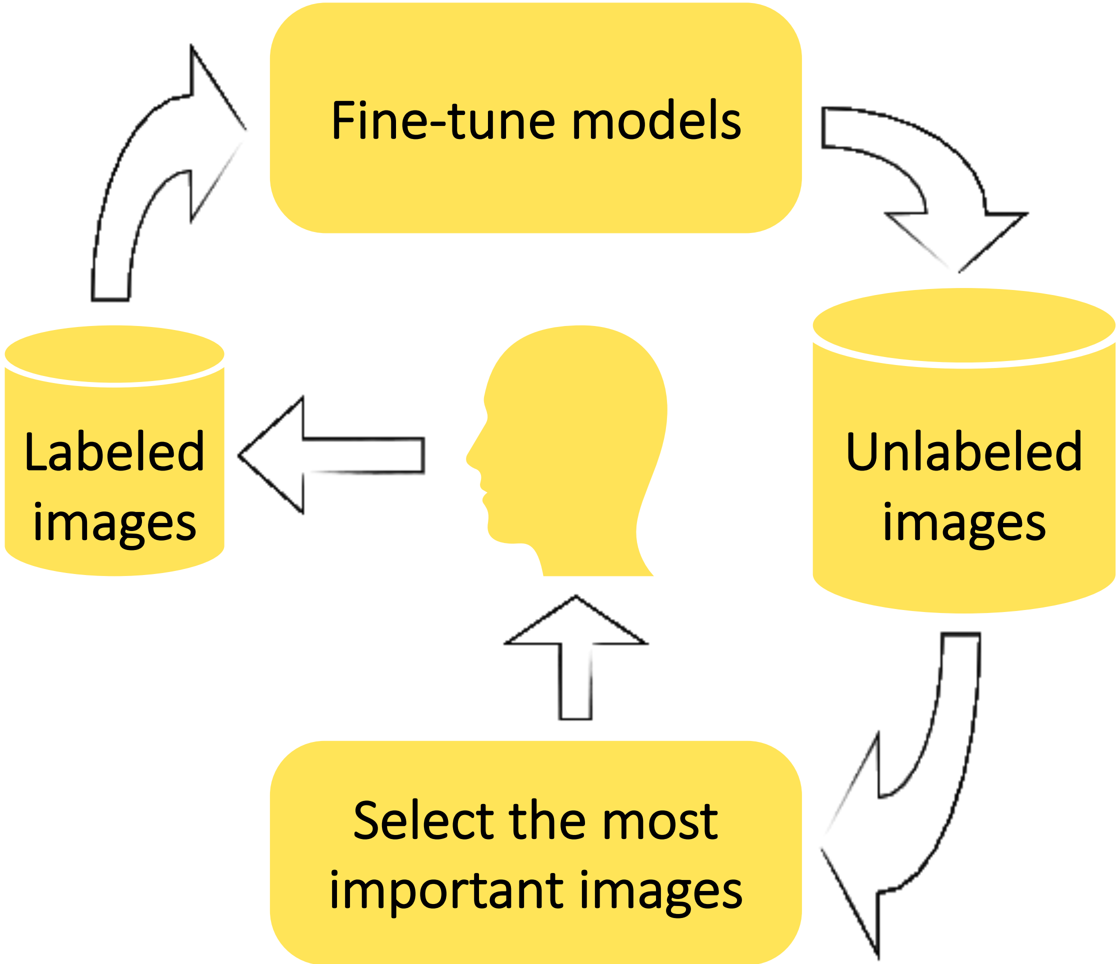

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

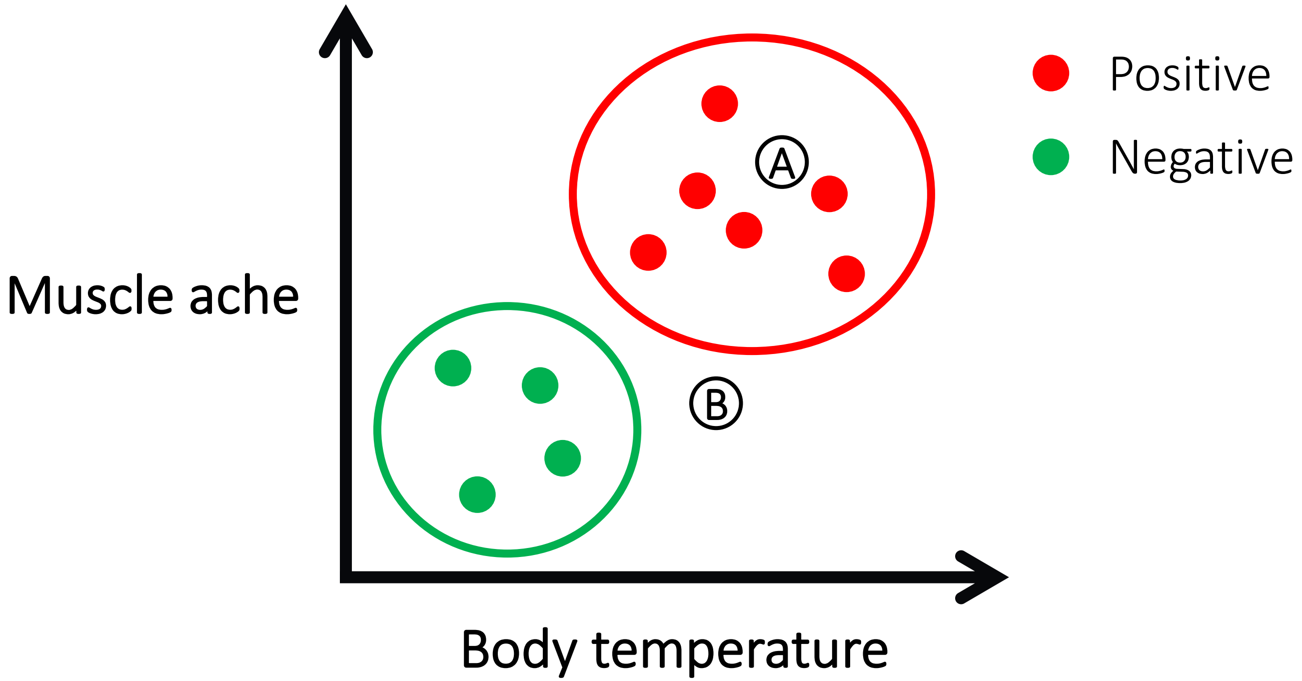

Select the most important samples

Given one dollar, which patient would you annotate, A or B?

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

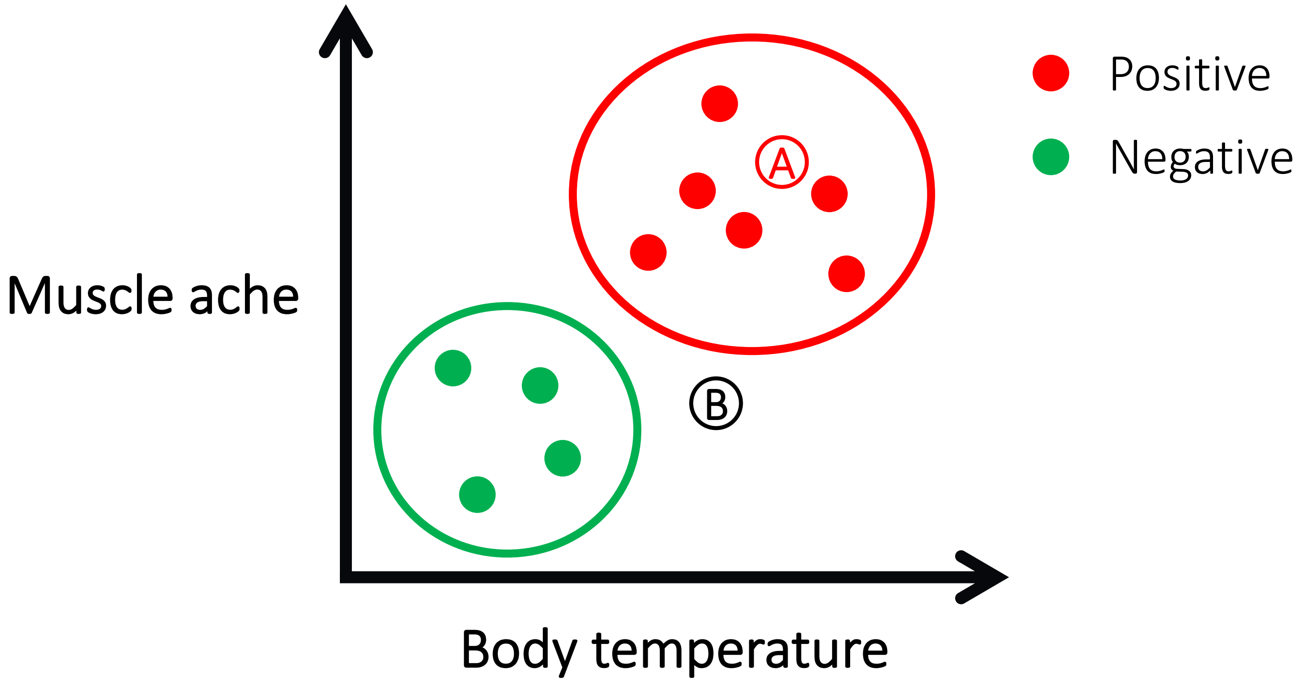

Select the most important samples

Given one dollar, which patient would you annotate, A or B?

Introduction

Significance

Aim #1

Aim #2

Aim #3

Summary

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

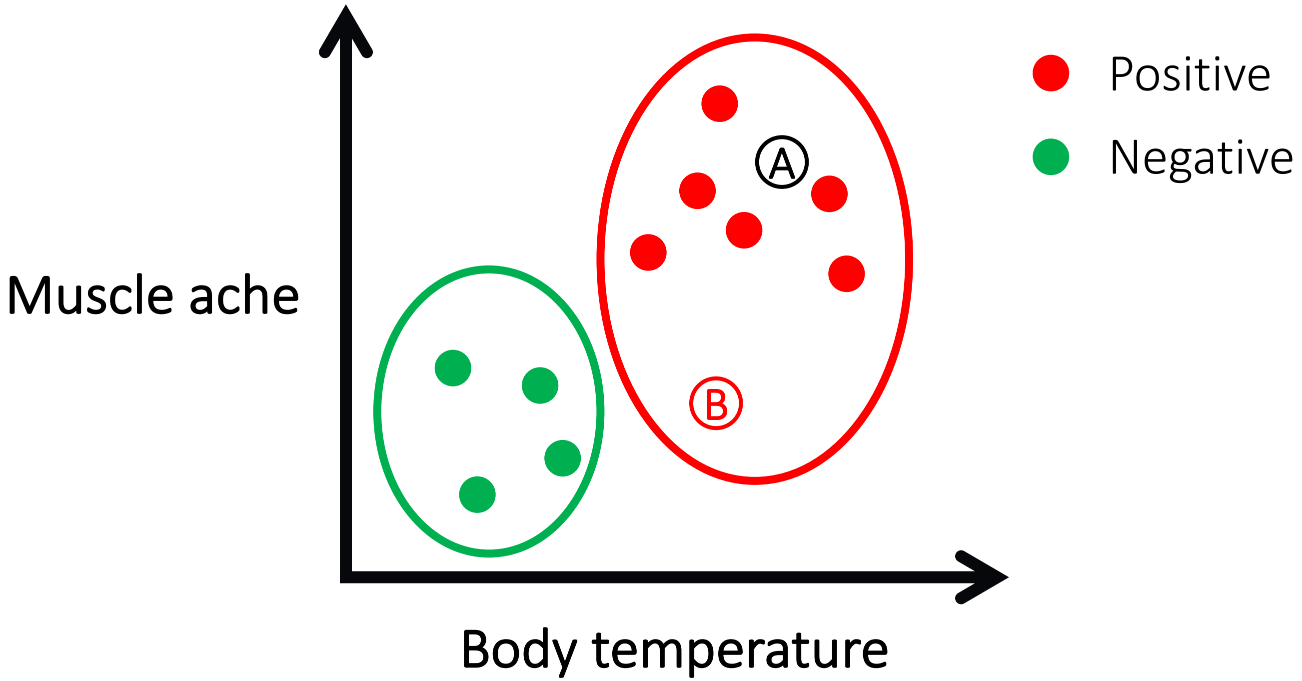

Select the most important samples

Given one dollar, which patient would you annotate, A or B?

Aim #1: Acquire necessary annotation efficiently from human experts

Approach: "Human-in-the-loop" active learning procedure

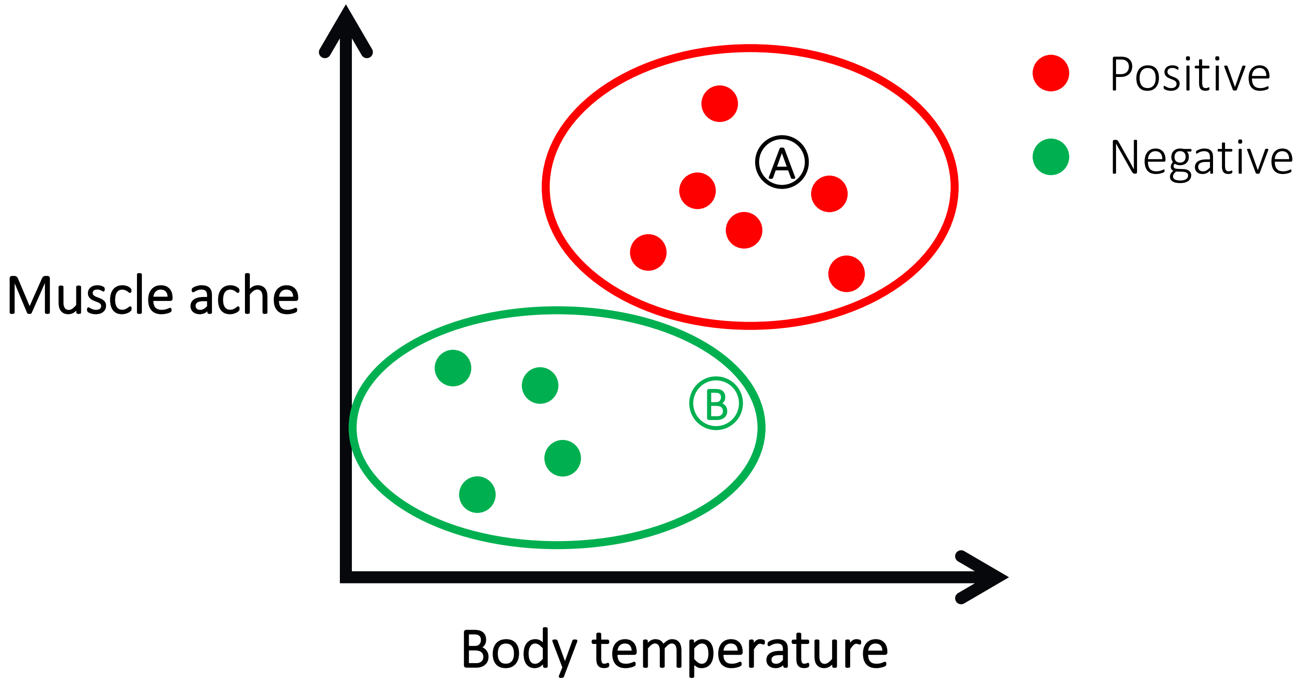

Select the most important samples

Given one dollar, which patient would you annotate, A or B?

Aim #1: Acquire necessary annotation efficiently from human experts

Hypothesis: Wisely selecting important samples can reduce annotation cost

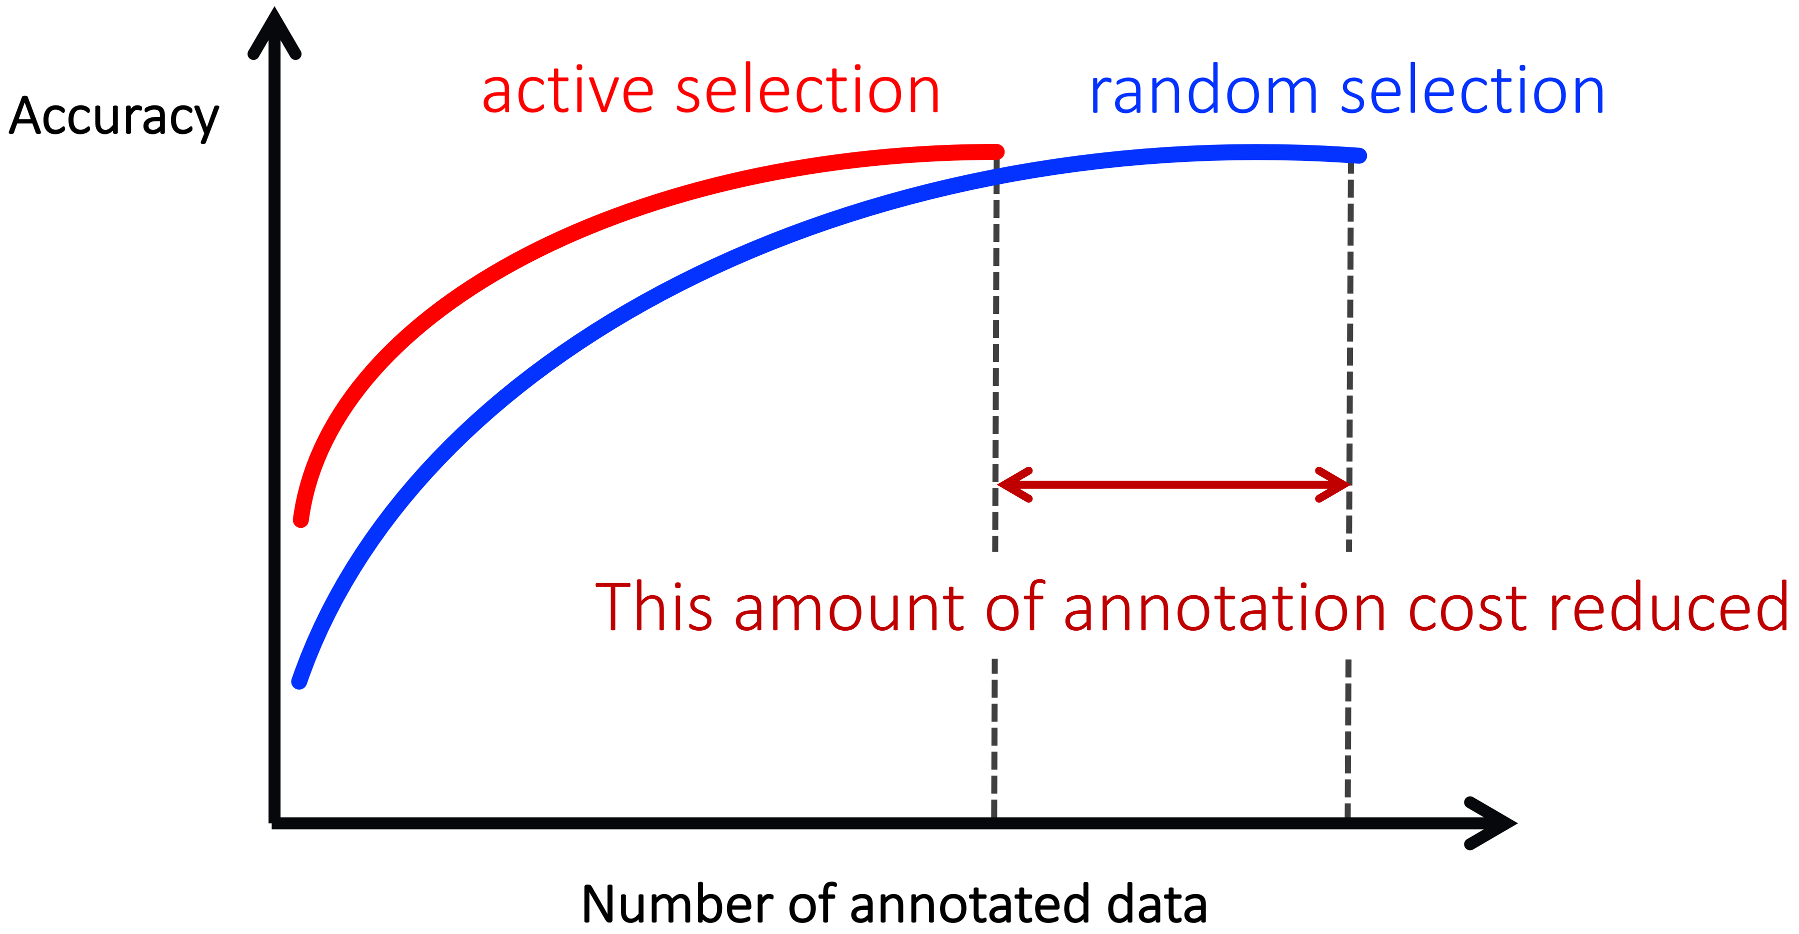

Aim #1: Acquire necessary annotation efficiently from human experts

Contribution: Reduce annotation cost by >60% compared to random selection

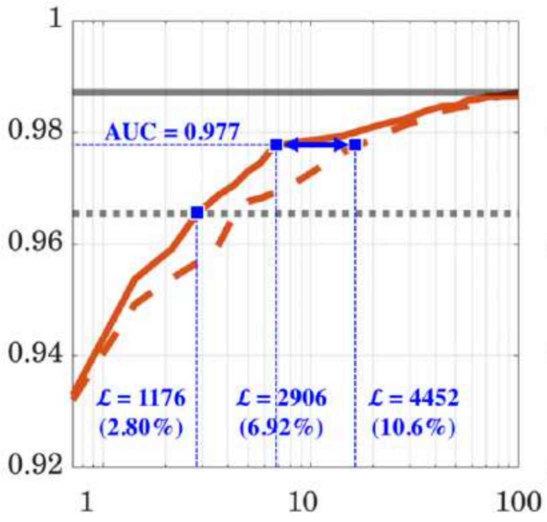

Places: Scene Classification

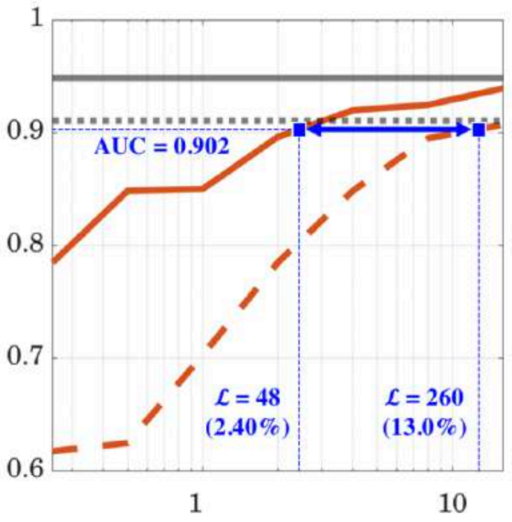

Colonoscopy Frame Classification

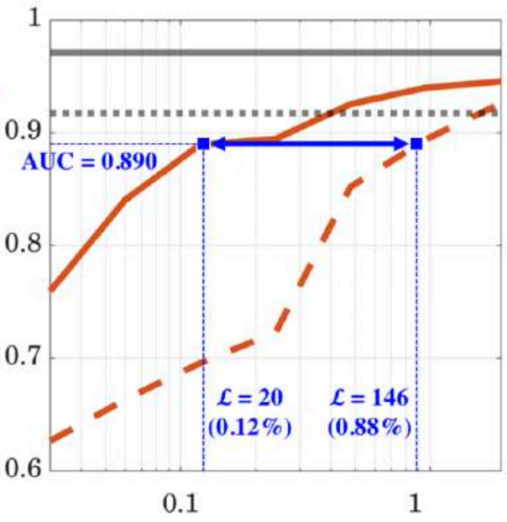

Polyp Detection

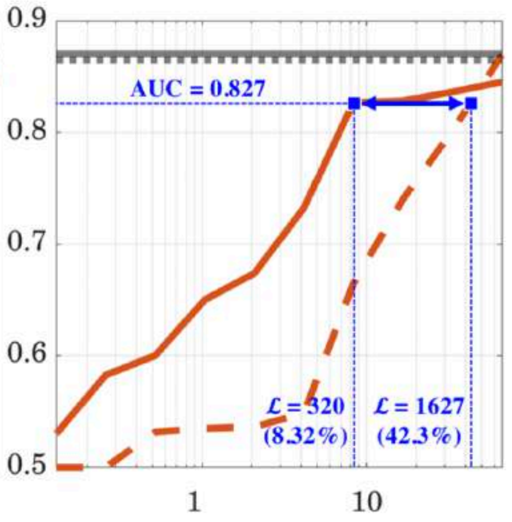

Pulmonary Embolism Detection

- Z. Zhou, *et al*. Integrating Active Learning and Transfer Learning for Carotid Intima-Media Thickness Video Interpretation. *Journal of Digital Imaging*, 2019. (IF=2.57)
- Z. Zhou, *et al*. Active Fine Tuning of Convolutional Neural Networks for Reducing Annotation Efforts. Submitted to *Medical Image Analysis*. (IF=8.88)
- Z. Zhou, *et al*. Fine-tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally. *CVPR'17*.

Aim #1: Acquire necessary annotation efficiently from human experts

Proposal: Iteratively suggest important samples at the patient-level

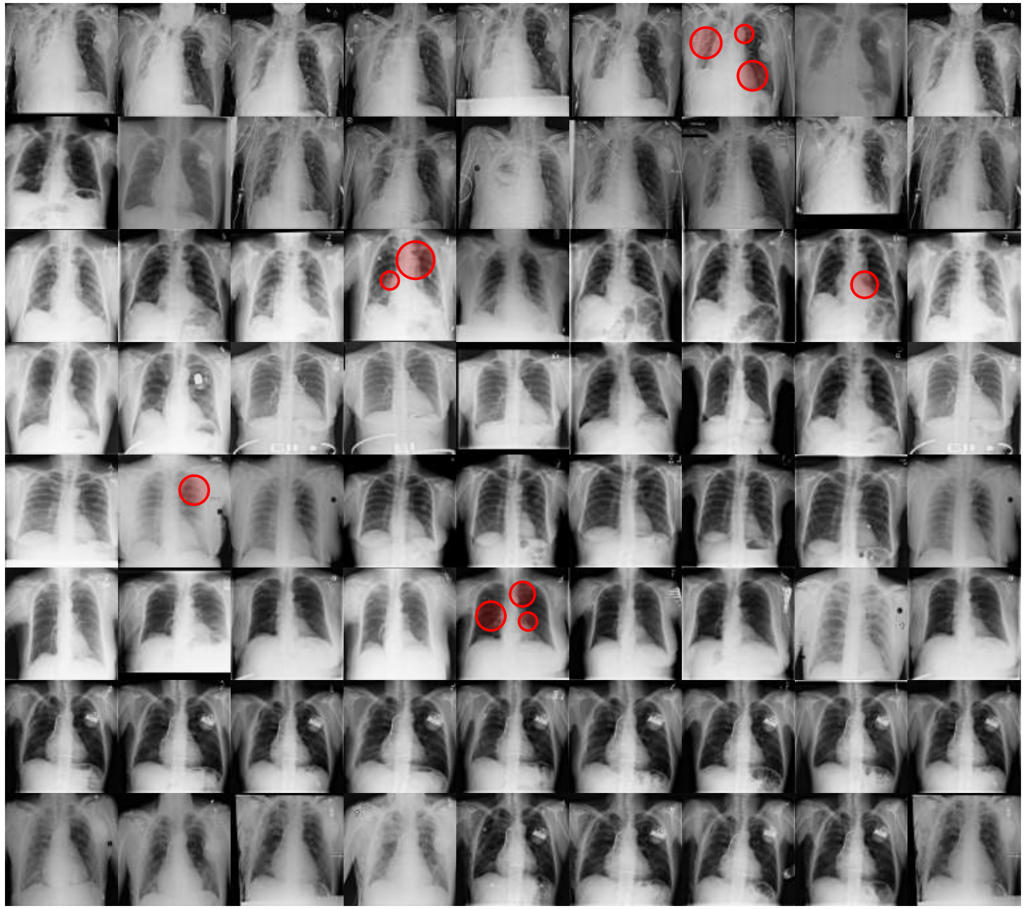

Lesion-level annotation

Drawbacks:

- Experts must annotate the same patient multiple times

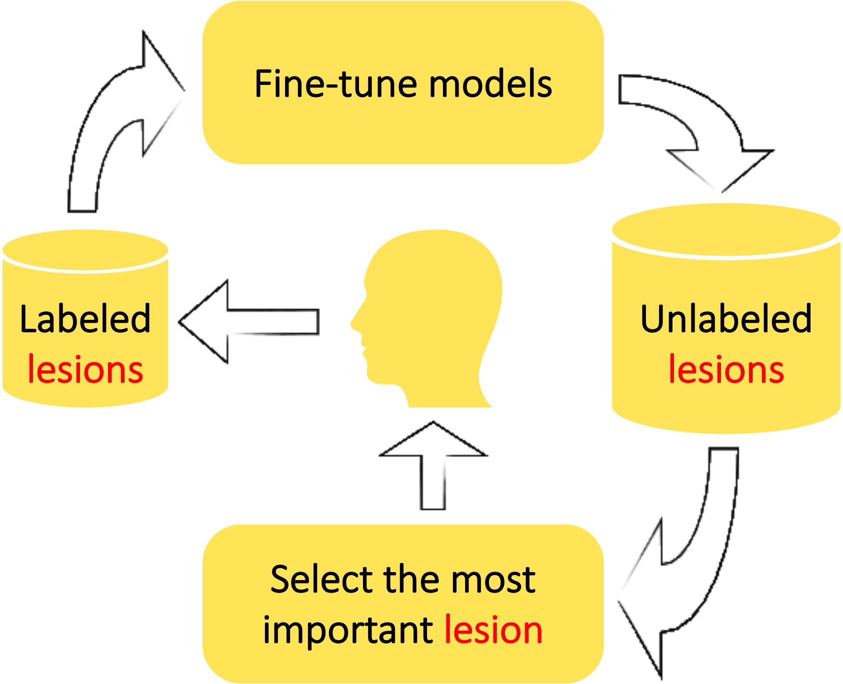

Aim #1: Acquire necessary annotation efficiently from human experts

Proposal: Iteratively suggest important samples at the patient-level

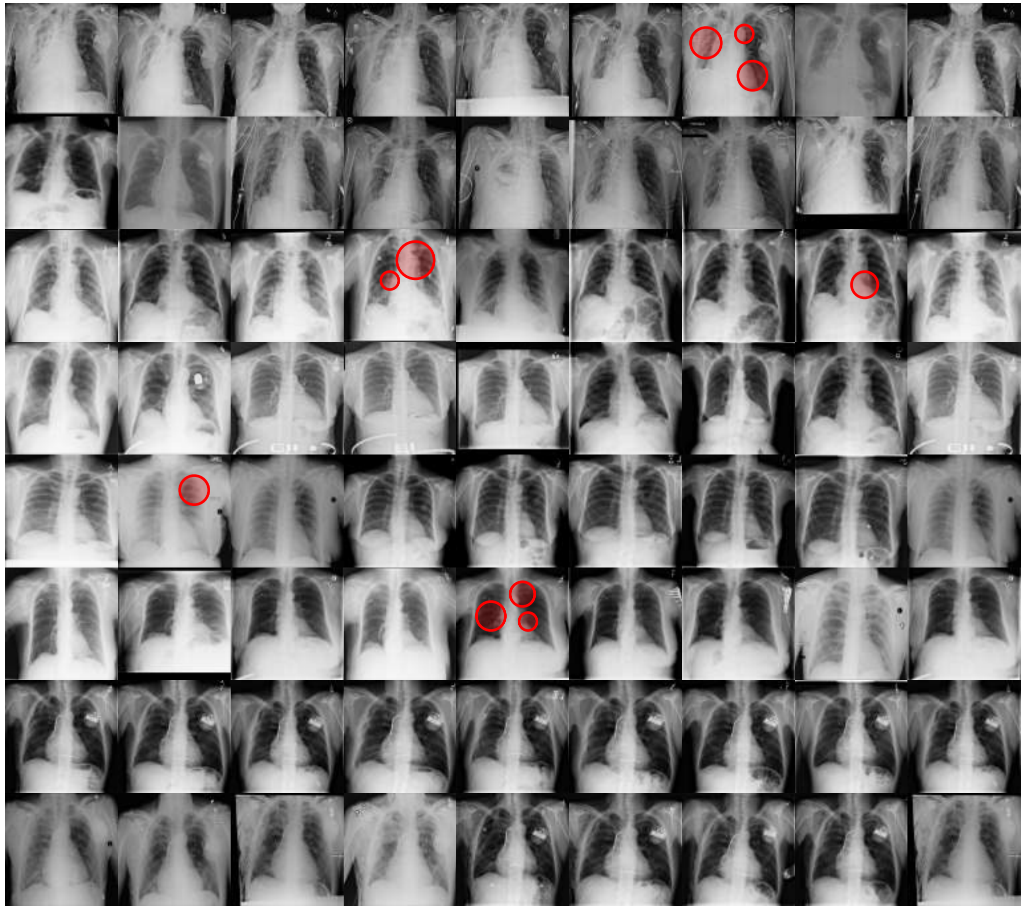
Lesion-level annotation

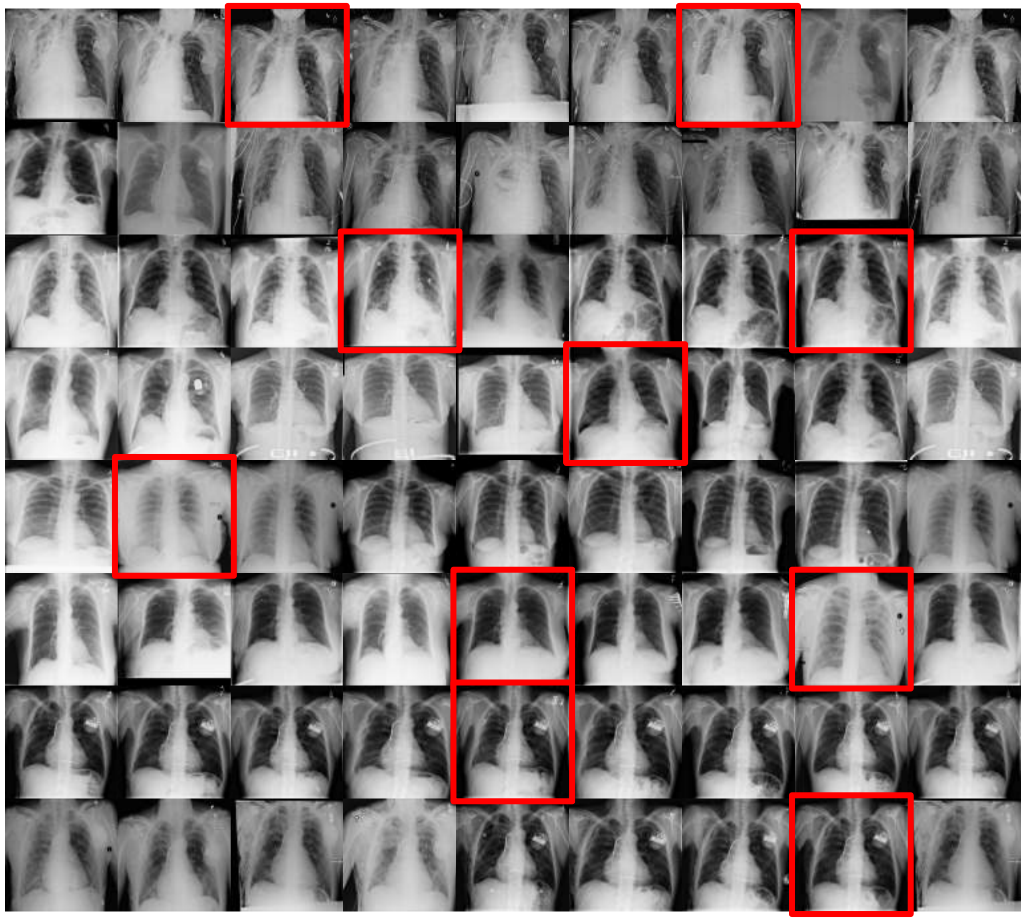
Patient-level annotation

Aim #2: Utilize existing annotation effectively from advanced architecture

Problem: Enhance the architecture for modeling 1,000 annotated images

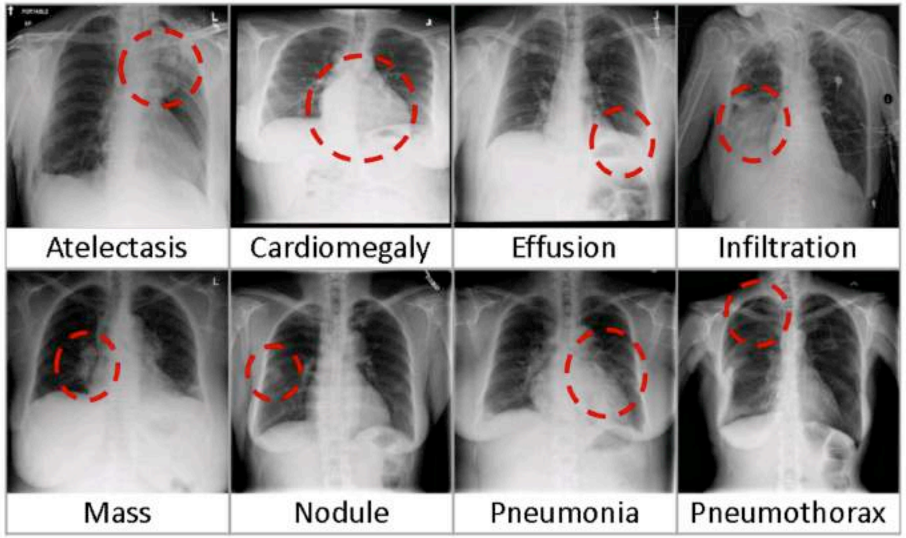

$ 1,000 annotation budget ☺

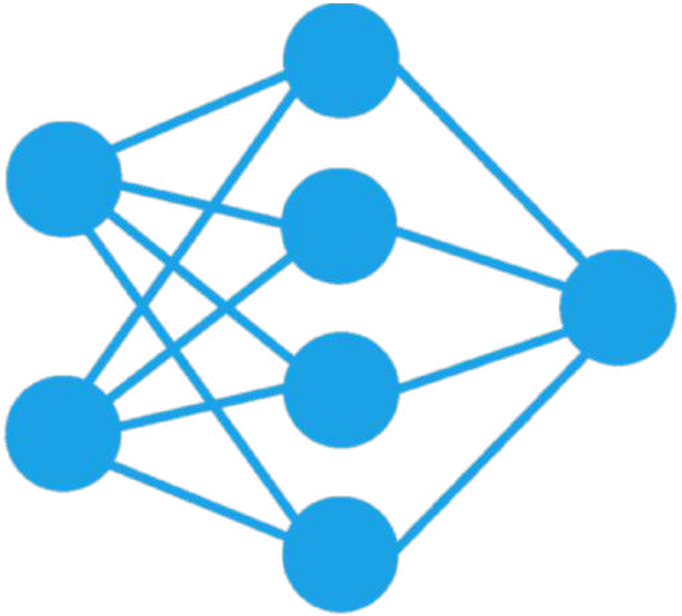

Aim #2: Utilize existing annotation effectively from advanced architecture

Segmentation: Partition an image into multiple segments to ease the analysis

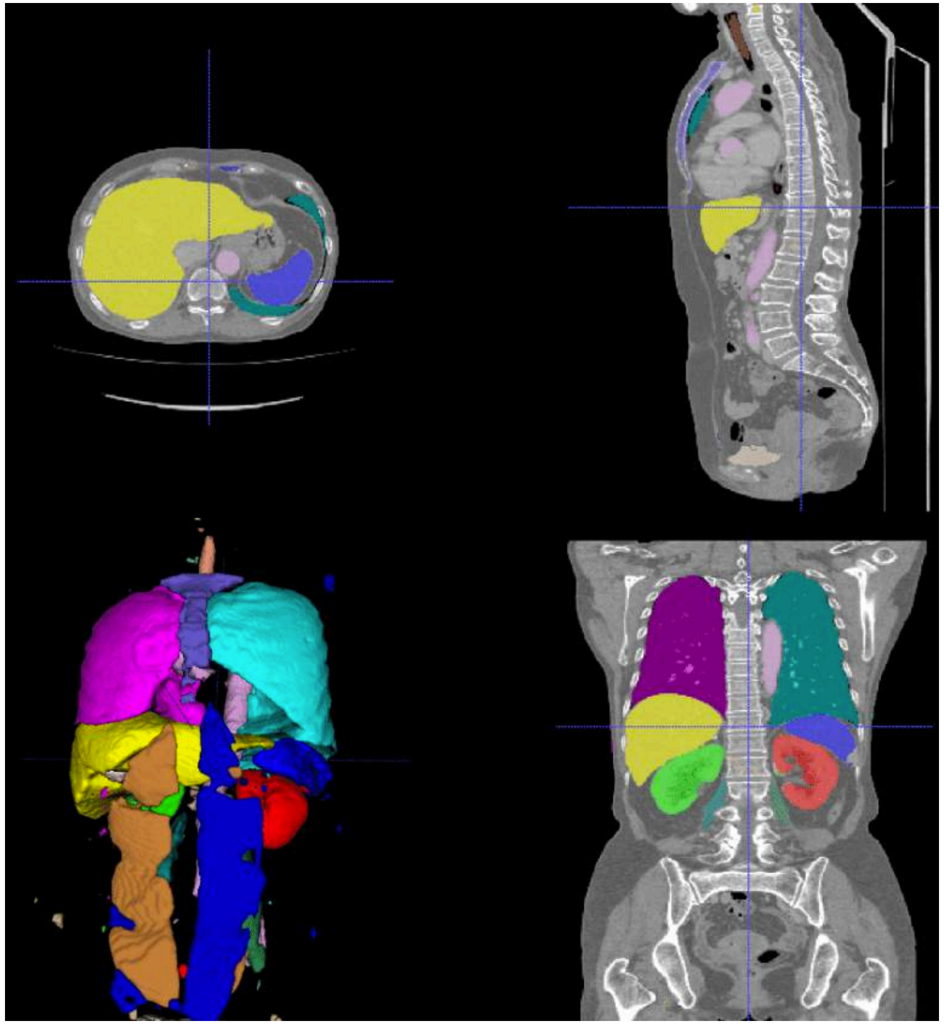

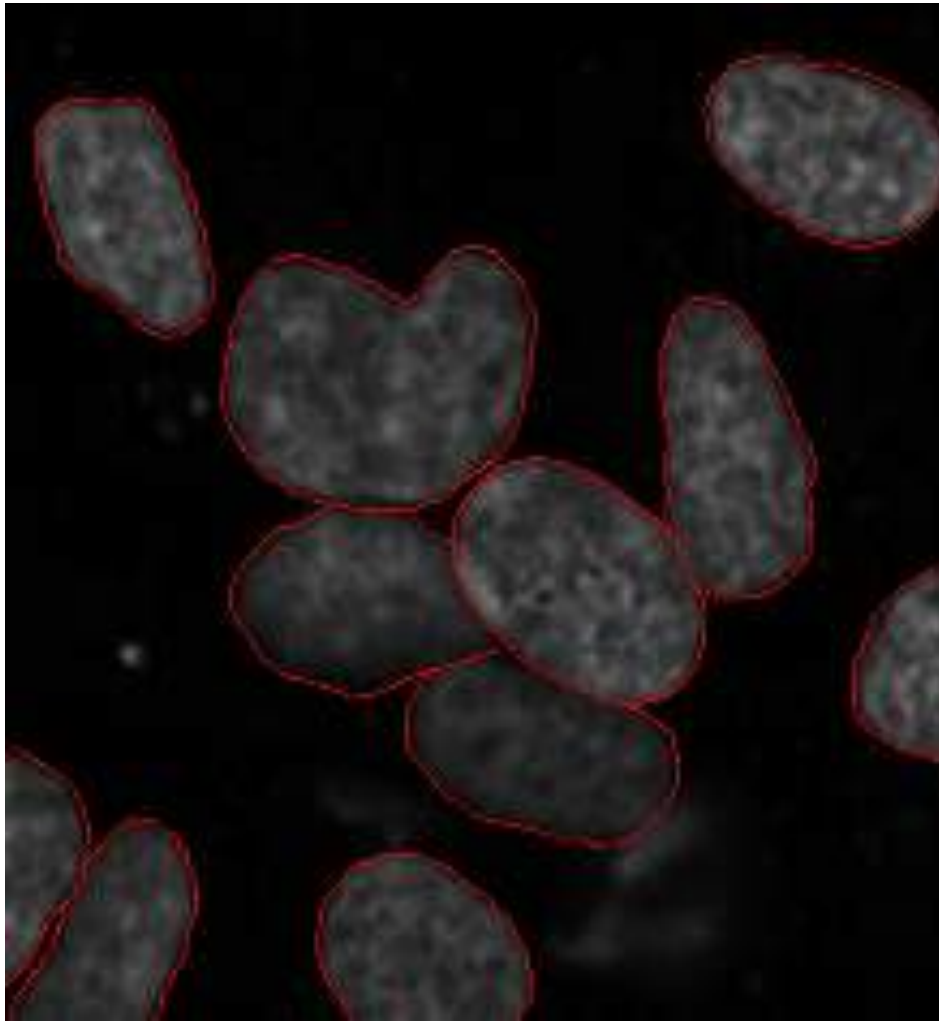

Aim #2: Utilize existing annotation effectively from advanced architecture

Segmentation: Partition an image into multiple segments to ease the analysis

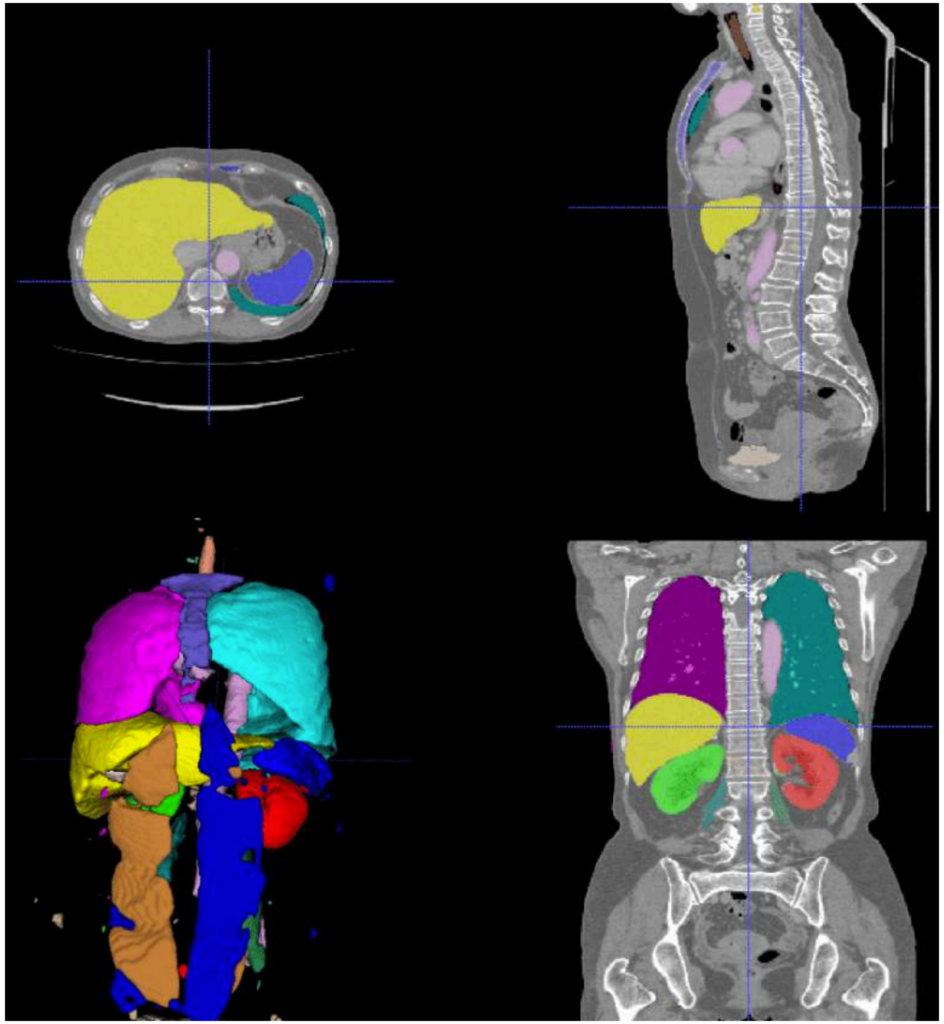

e.g., liver segmentation

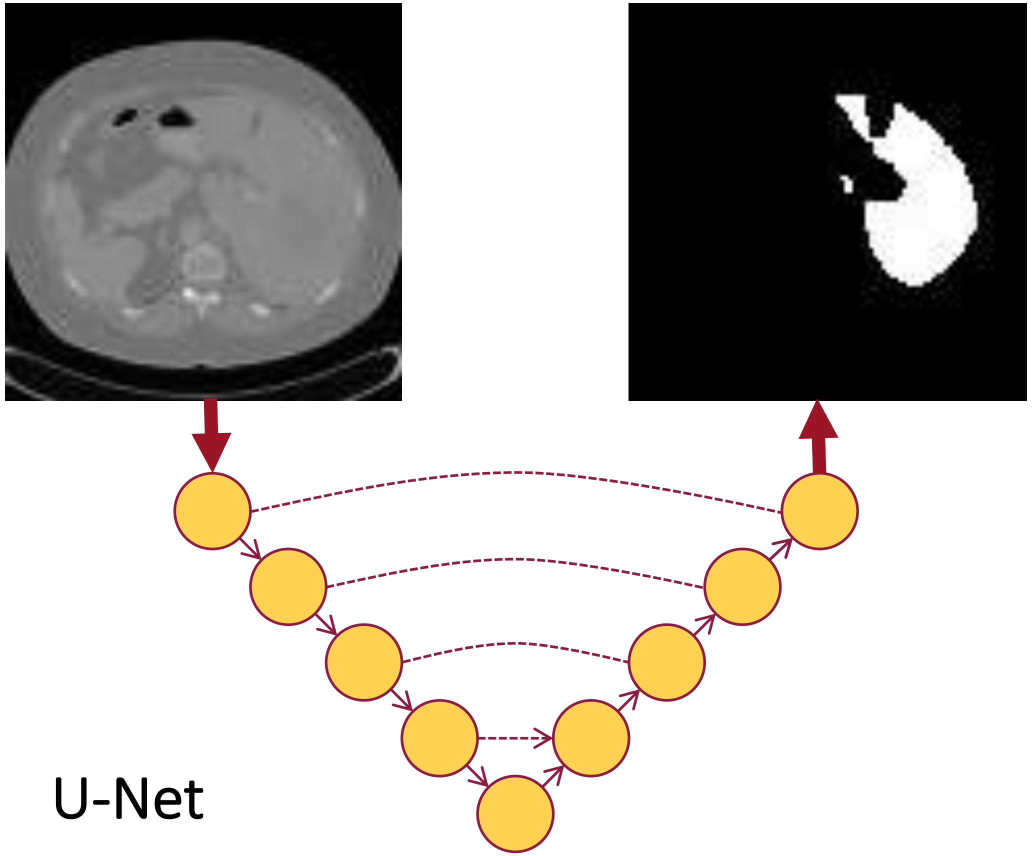

Aim #2: Utilize existing annotation effectively from advanced architecture

Hypothesis: Multi-scale feature aggregation leads to powerful models

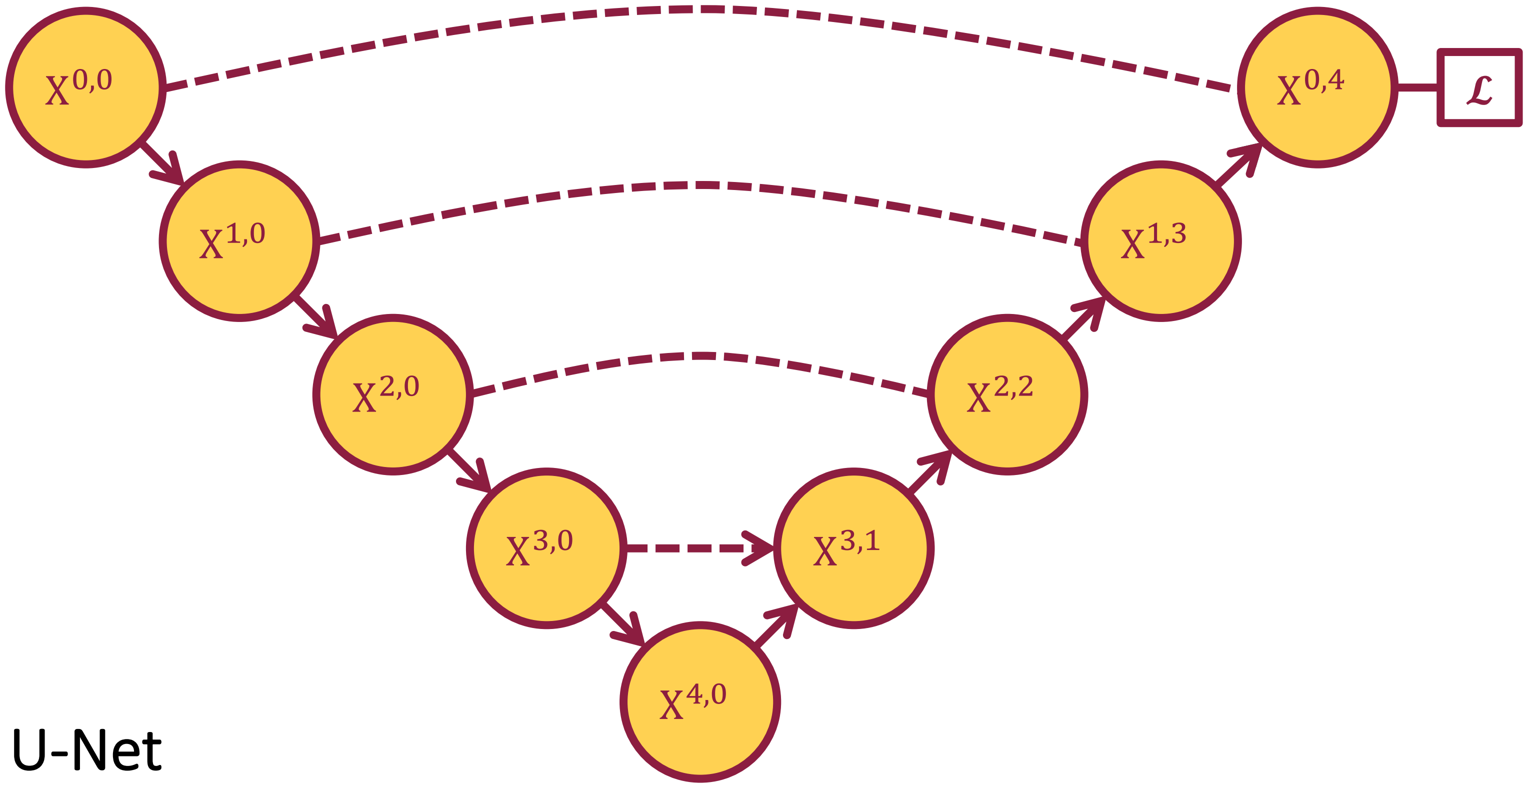

U-Net

[Olaf Ronneberger, *et al*., 2015]

Aim #2: Utilize existing annotation effectively from advanced architecture

Approach: Redesigned skip connections aggregate multi-scale features

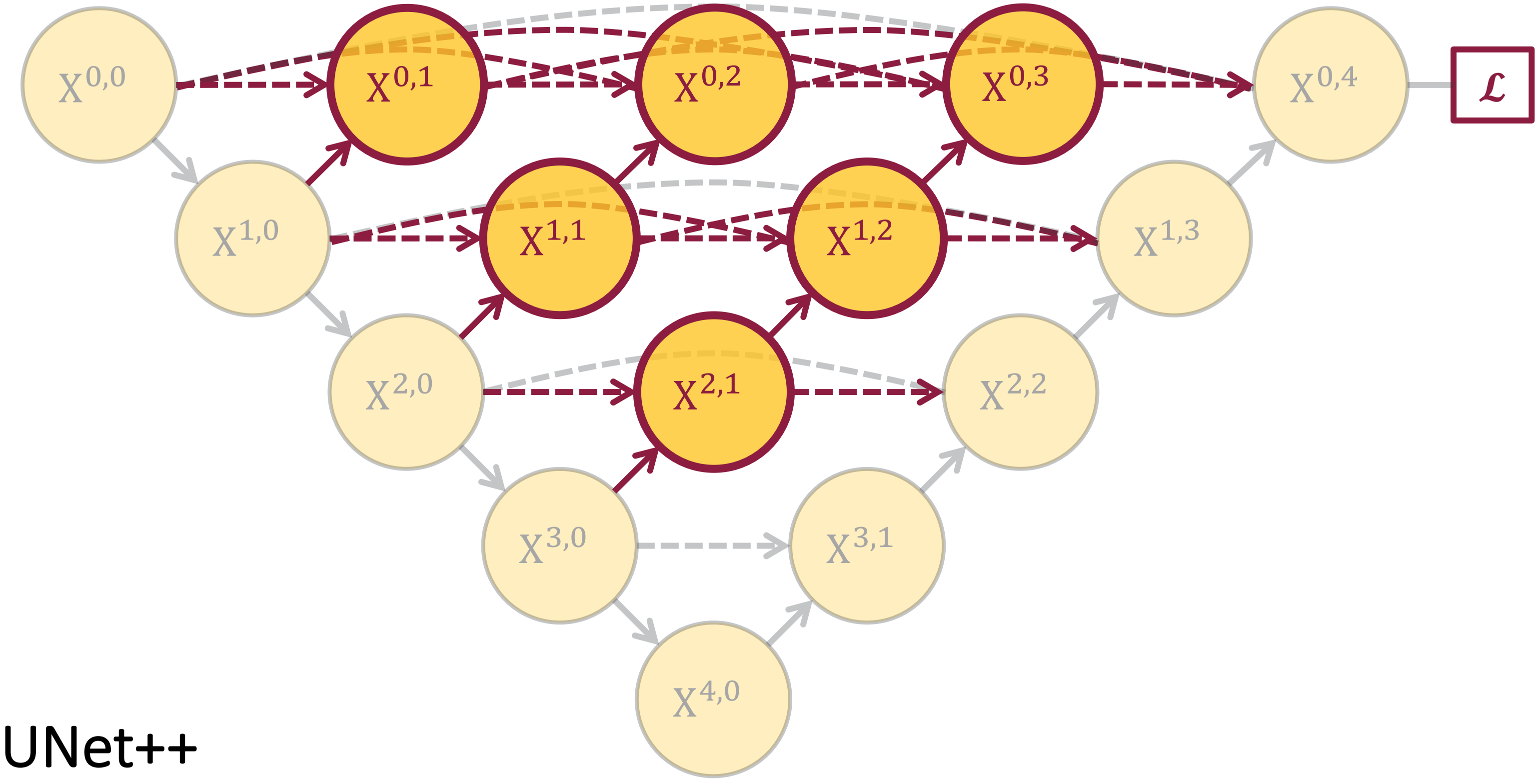

UNet++

[Zongwei Zhou, *et al.*, 2018]

Aim #2: Utilize existing annotation effectively from advanced architecture

Approach: Deep supervision enables a higher segmentation accuracy

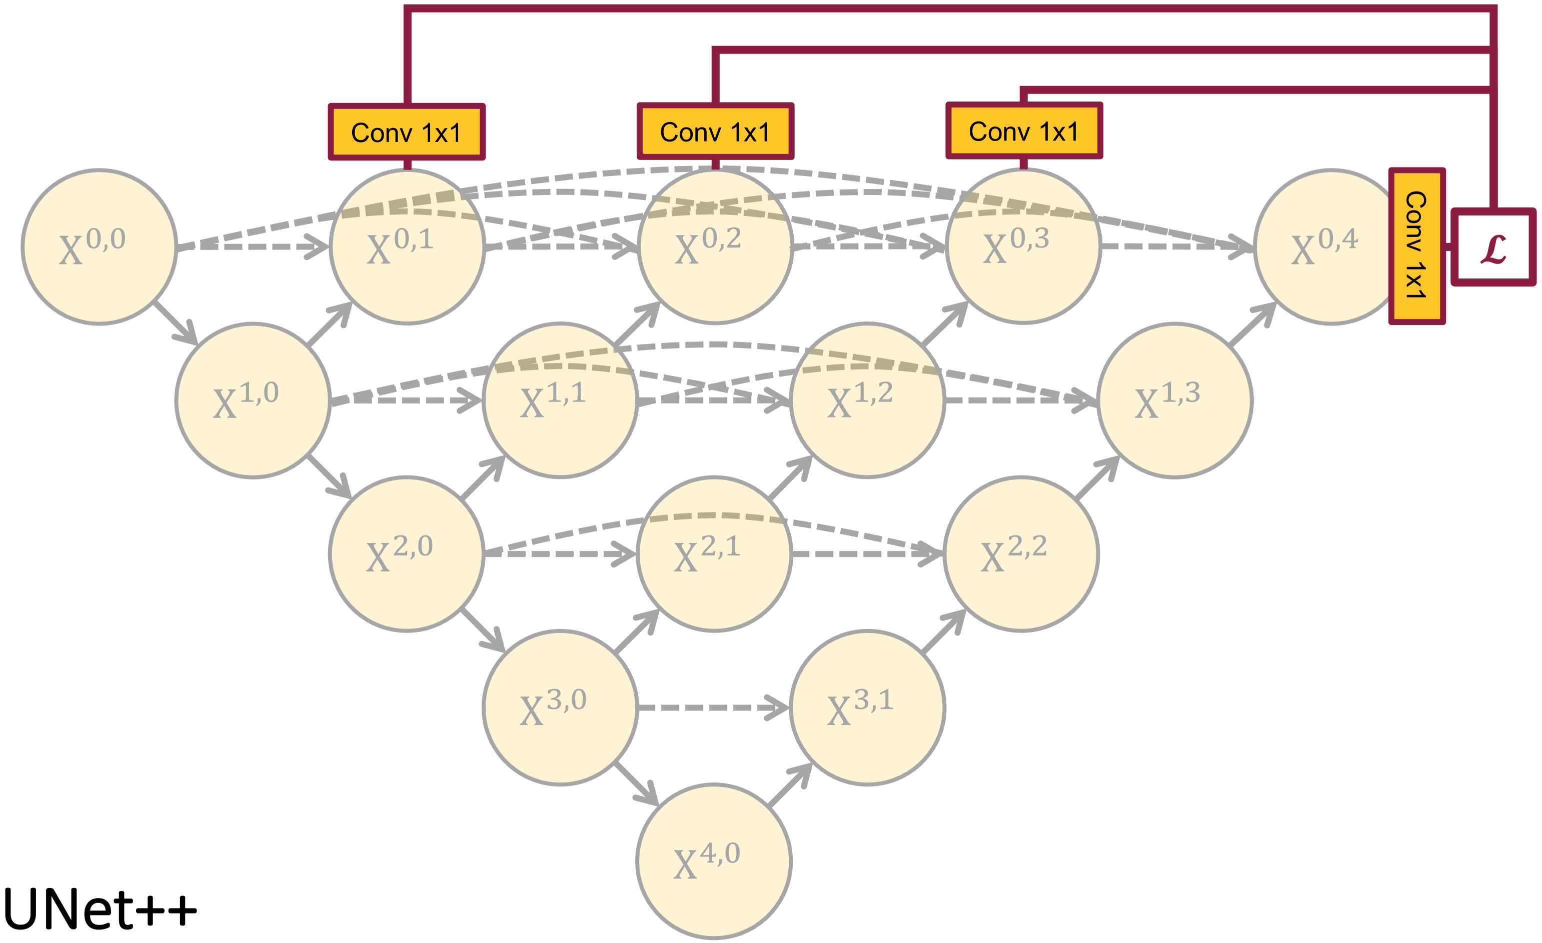

UNet++

[Zongwei Zhou, *et al.*, 2018]

Aim #2: Utilize existing annotation effectively from advanced architecture

Contribution: UNet++ significantly improves disease/organ segmentation

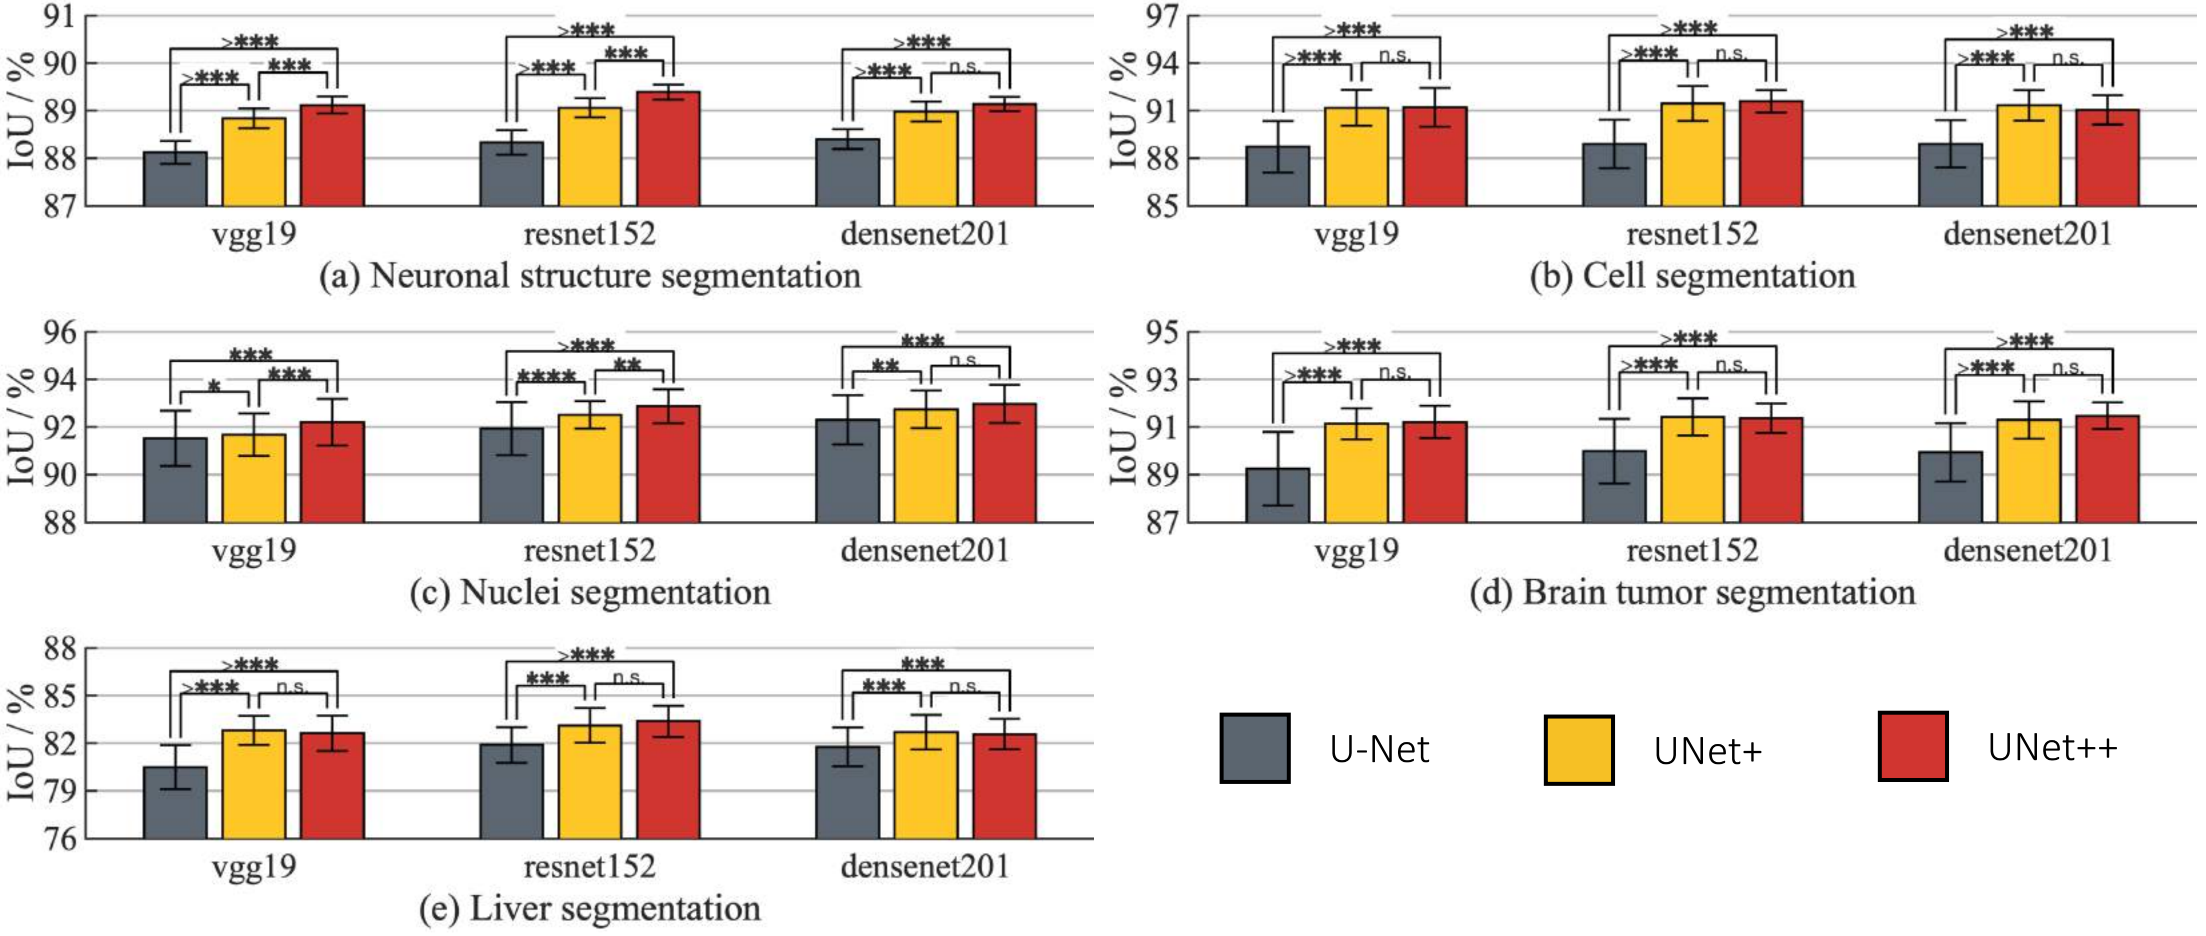

- Z. Zhou, *et al*. UNet++: Redesigning Skip Connections to Exploit Multi-Resolution Features in Image Segmentation. IEEE Transactions on Medical Imaging, 2020. (IF=7.82)
- Z. Zhou, et al. UNet++: A Nested U-Net Architecture for Medical Image Segmentation. *DLMIA'18*. (Oral)

Aim #2: Utilize existing annotation effectively from advanced architecture

Approach: Deep supervision allows model pruning

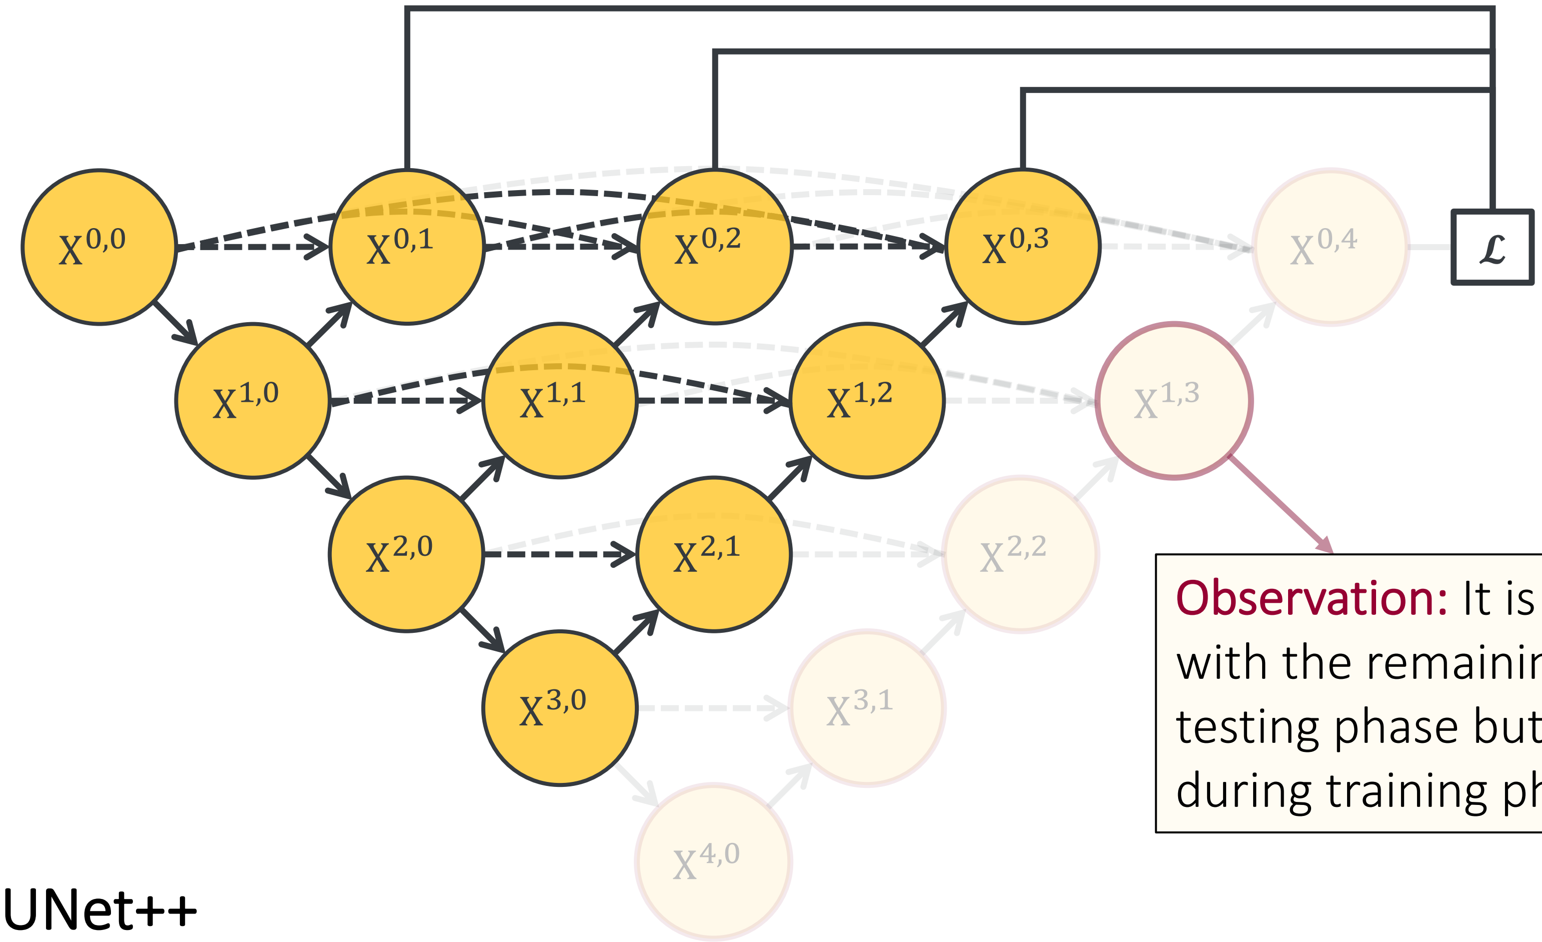

UNet++

Observation: It is independent with the remaining parts during testing phase but contributes during training phase.

Aim #2: Utilize existing annotation effectively from advanced architecture

Contribution: Pruned UNet++ accelerates a computer reading medical images

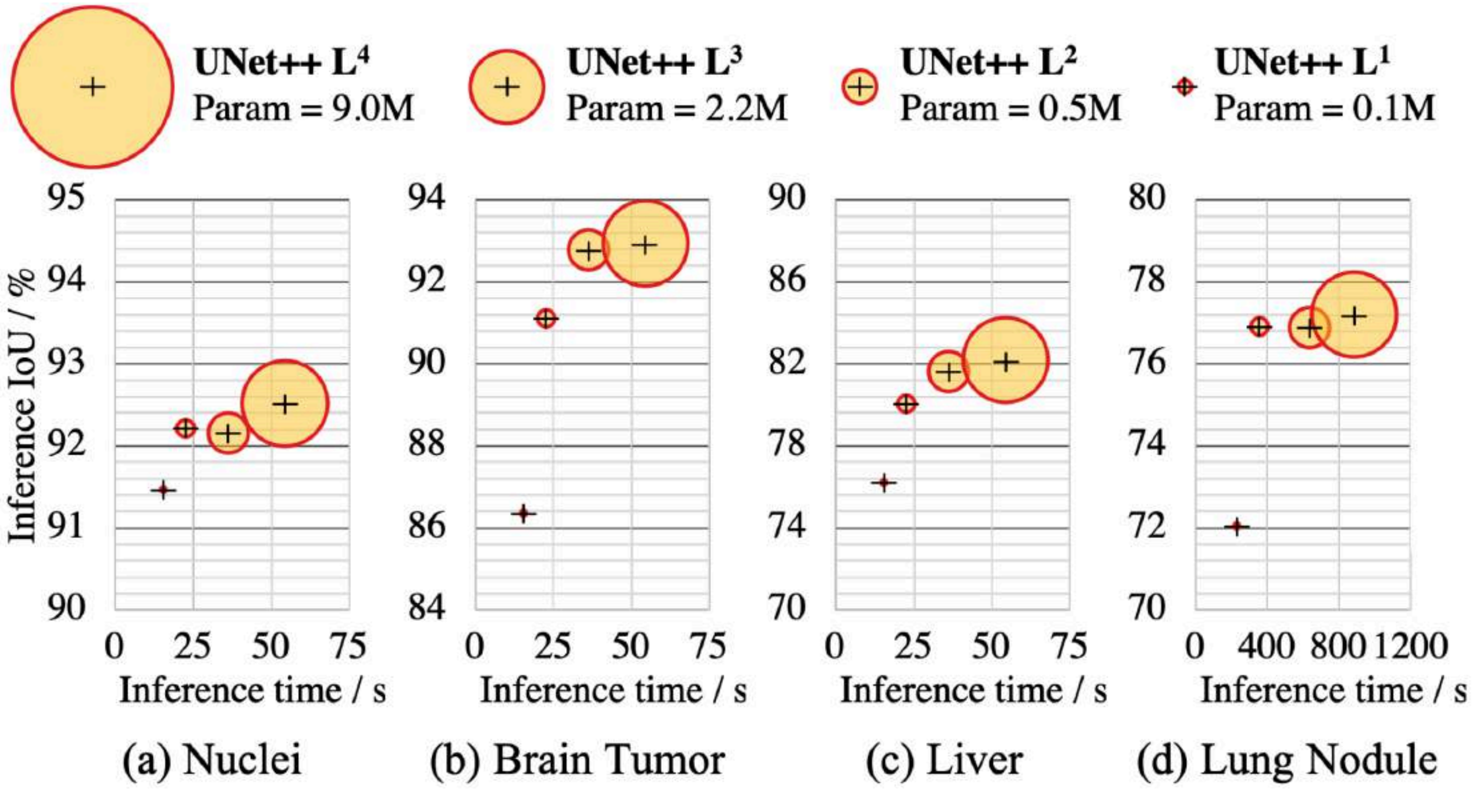

- Z. Zhou, *et al*. UNet++: Redesigning Skip Connections to Exploit Multi-Resolution Features in Image Segmentation. IEEE Transactions on Medical Imaging, 2020. (IF=7.82)
- Z. Zhou, et al. UNet++: A Nested U-Net Architecture for Medical Image Segmentation. *DLMIA'18*. (Oral)

Introduction

Significance

Aim #1

Aim #2

Aim #3

Summary

Aim #2: Utilize existing annotation effectively from advanced architecture

Proposal: Optimize active learning by leveraging unique architectural design

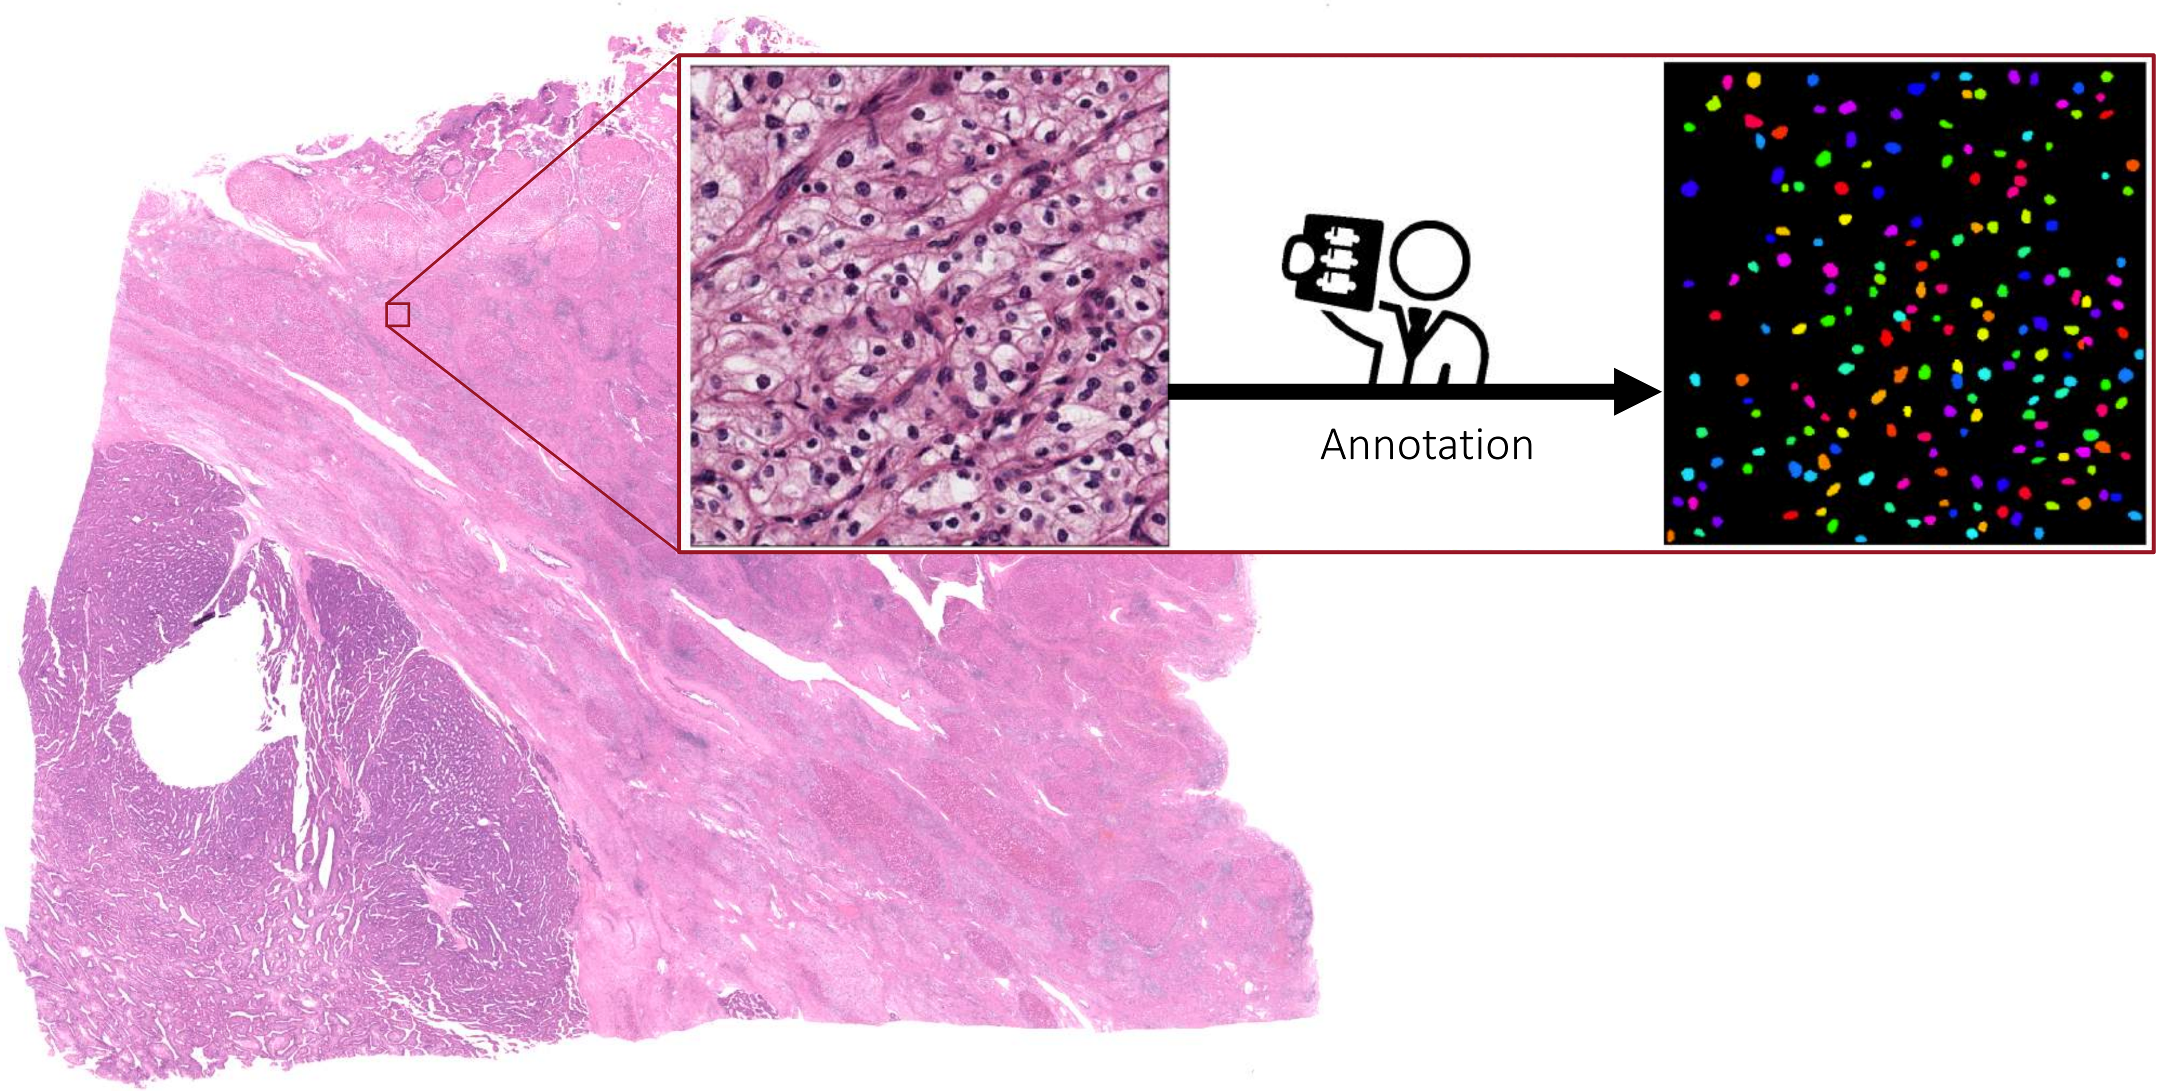

Aim #2: Utilize existing annotation effectively from advanced architecture

Proposal: Optimize active learning by leveraging unique architectural design

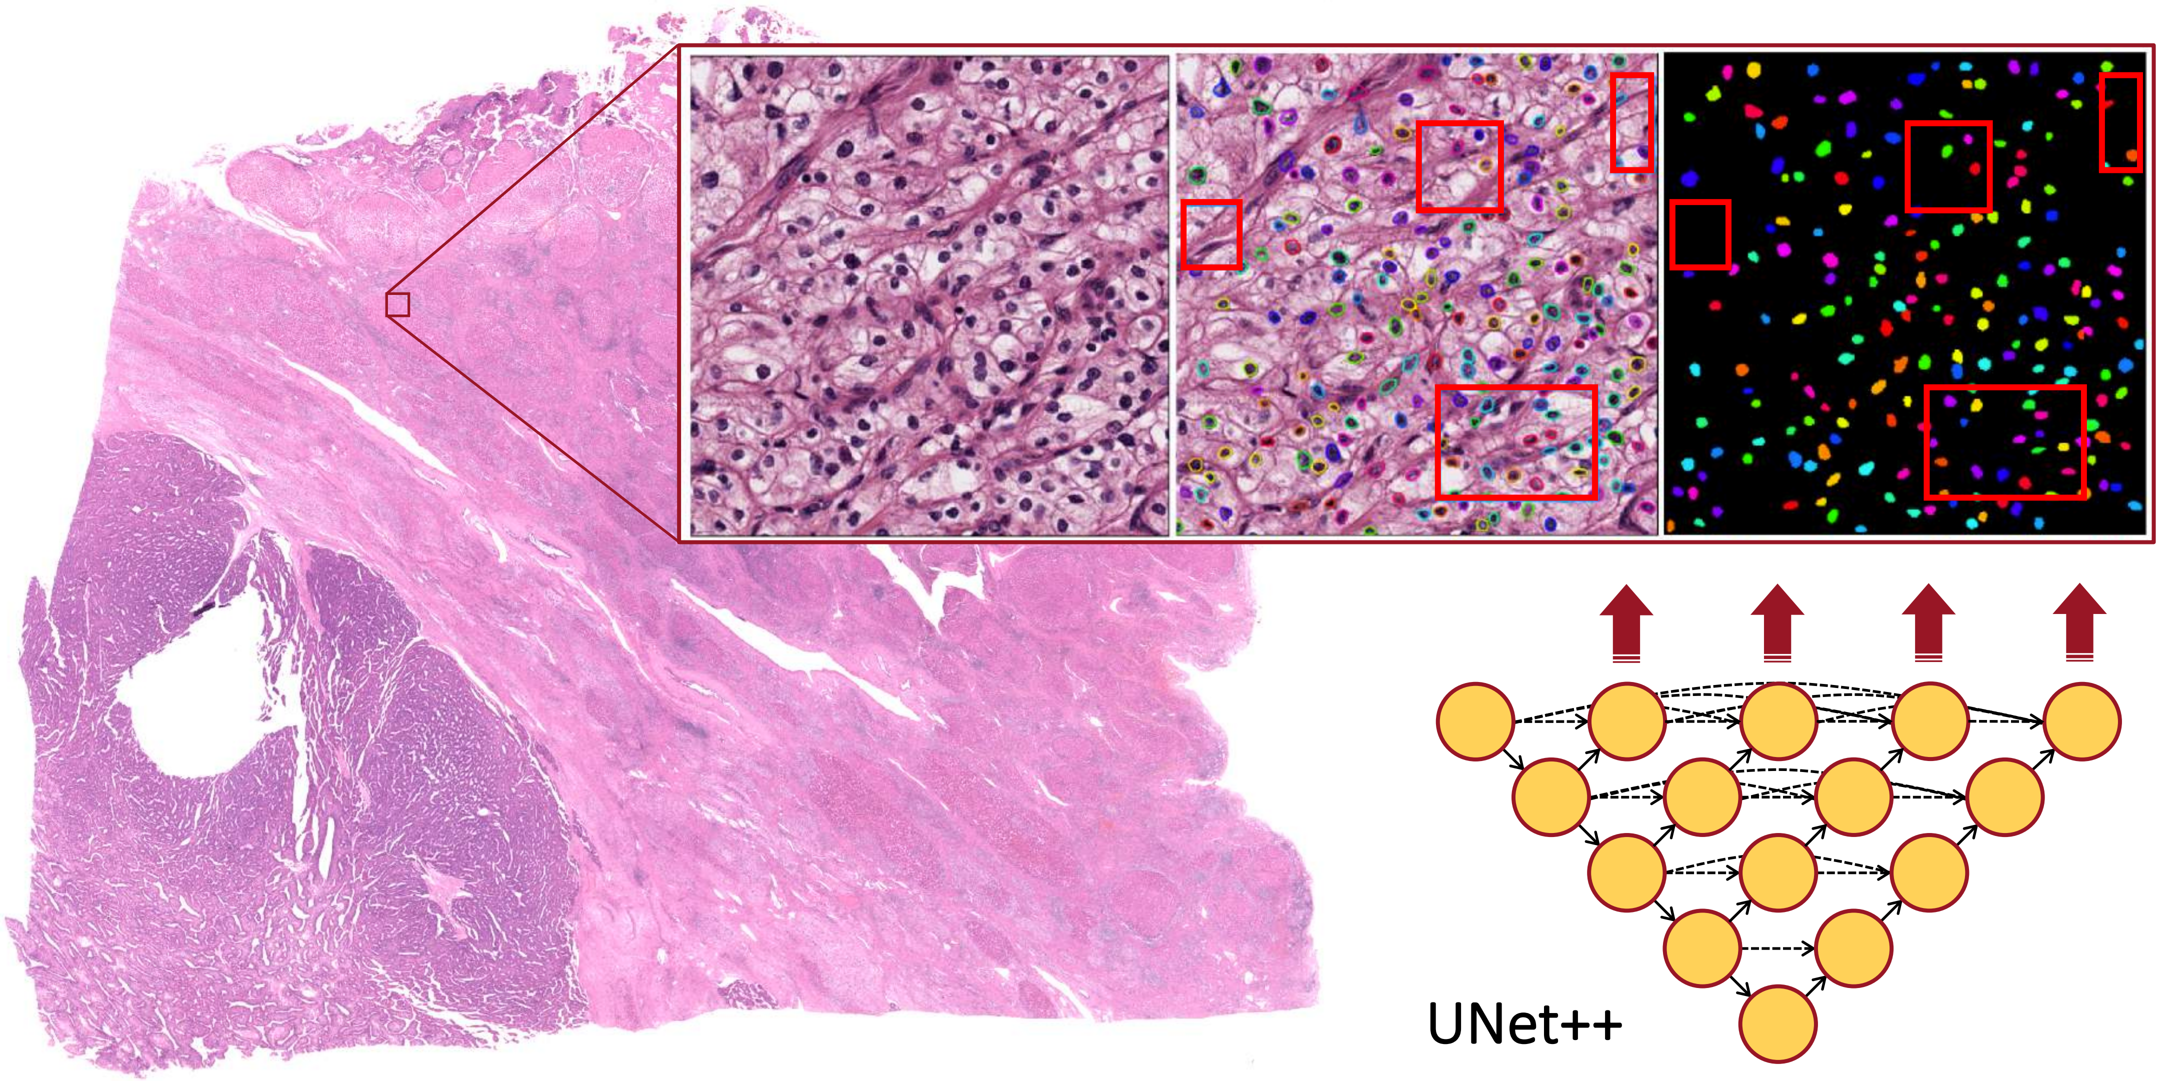

Aim #3: Extract generic knowledge directly from unannotated images

Problem: Utilize 1,000,000 images without systematic annotation

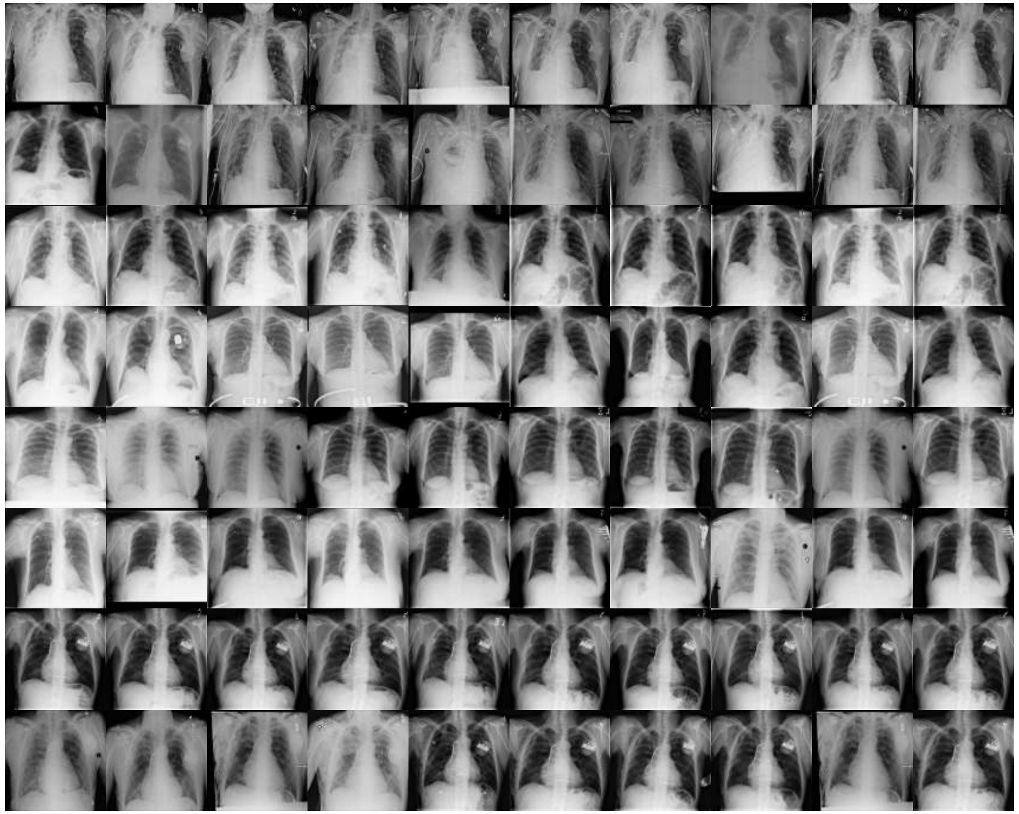

$ 1,000,000 annotation cost ☹

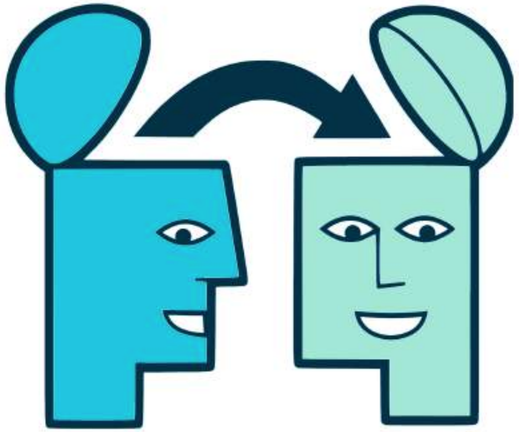

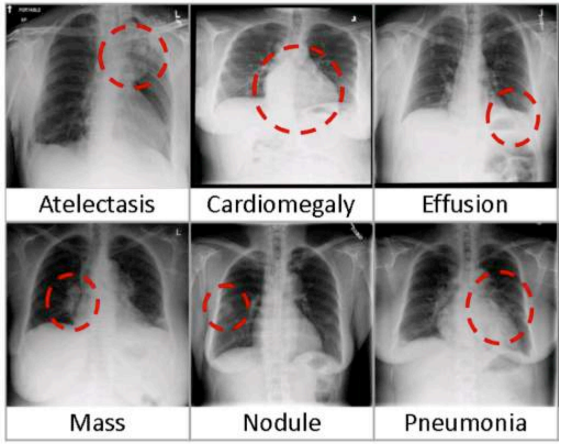

$ 100 annotation budget ☺

Aim #3: Extract generic knowledge directly from unannotated images

Hypothesis: Generic models can be built upon consistent, recurrent anatomy

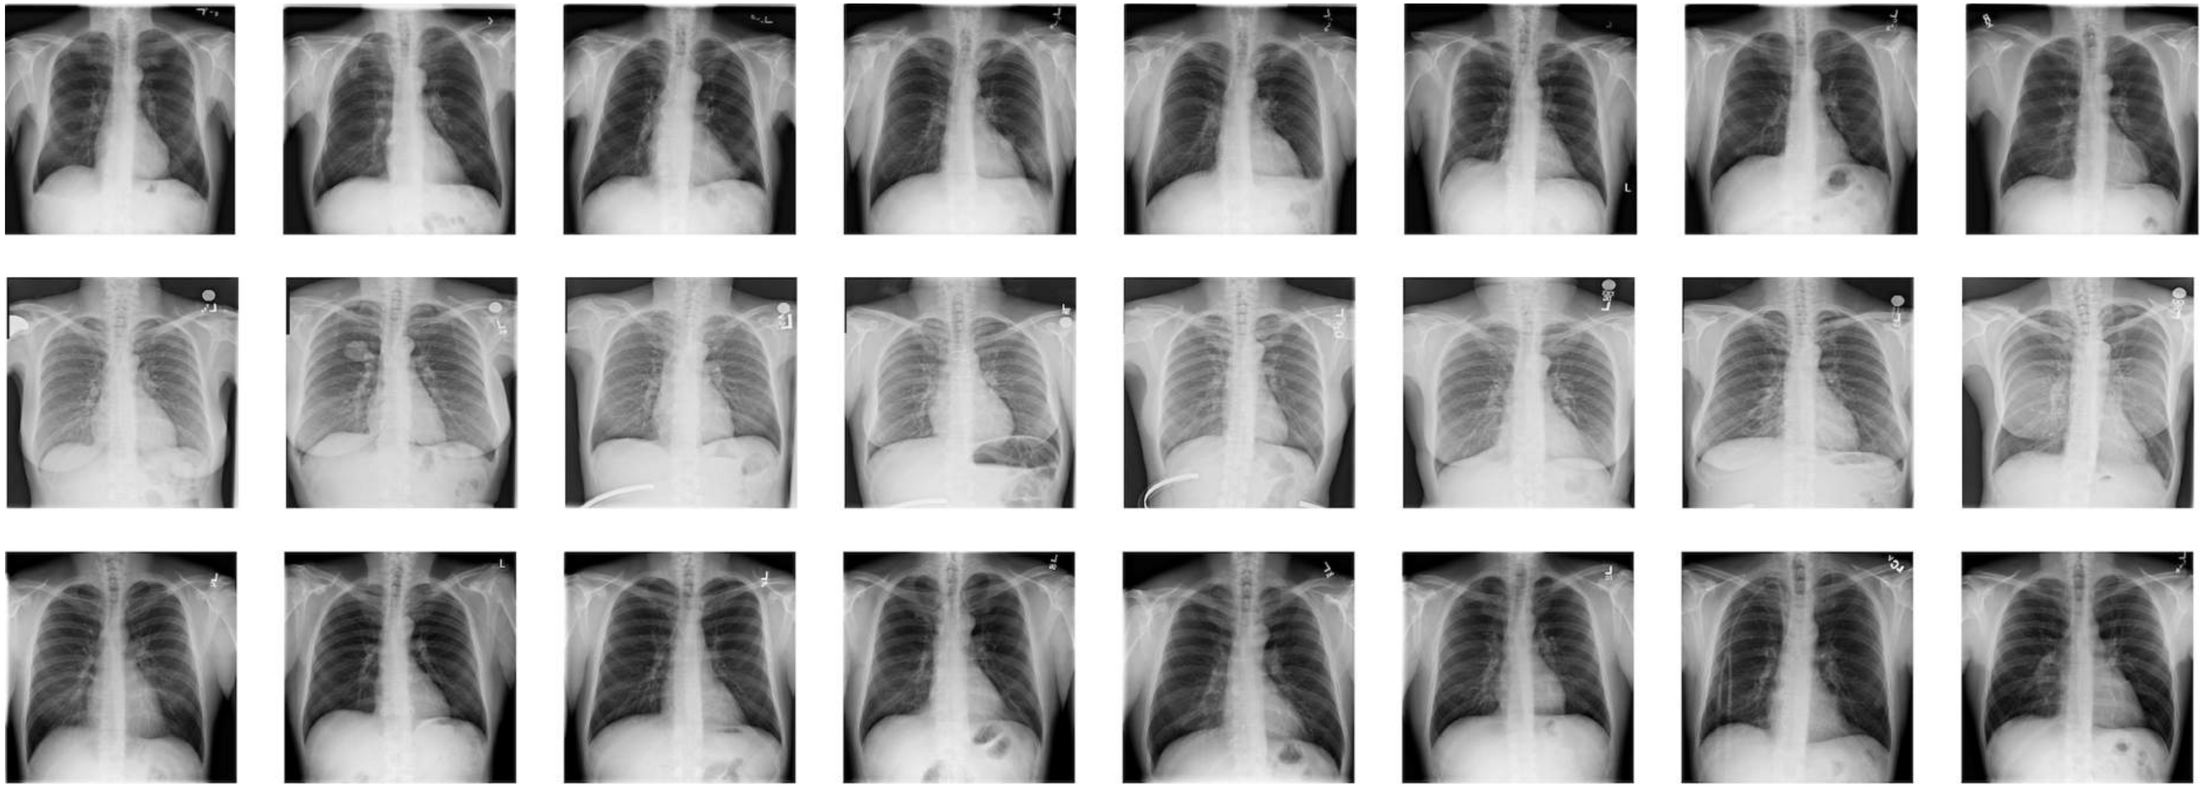

Aim #3: Extract generic knowledge directly from unannotated images

Approach: Image restoration task helps model learn image representation

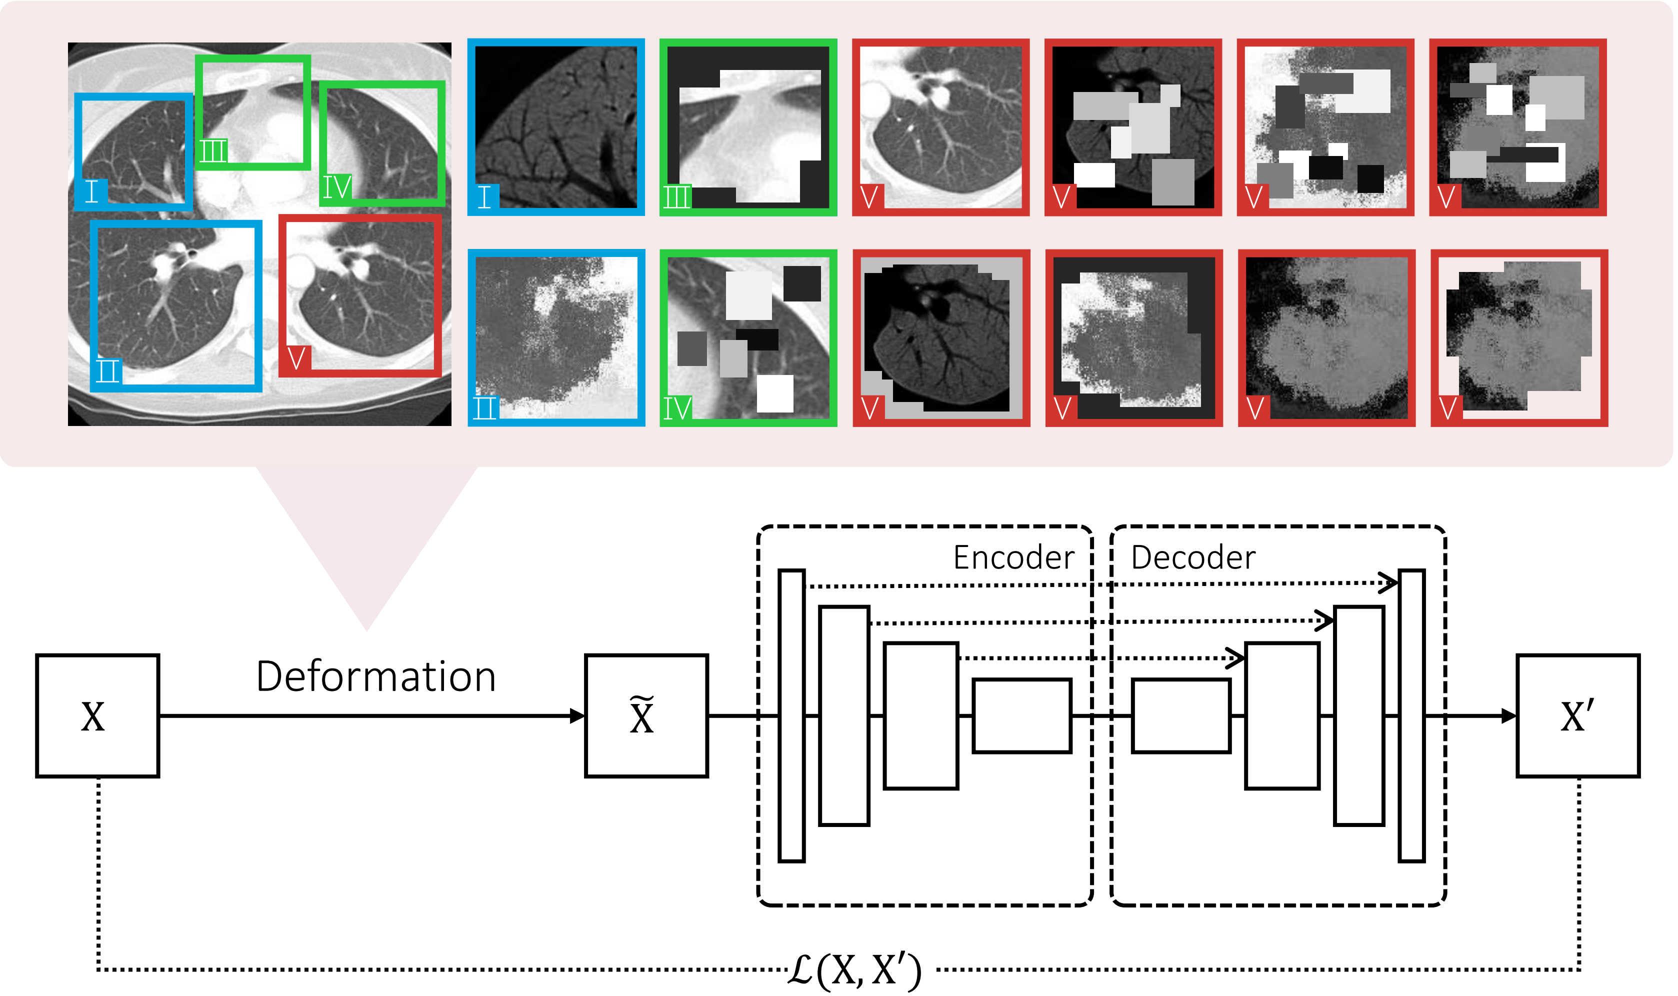

Aim #3: Extract generic knowledge directly from unannotated images

Approach: Learning from multiple perspectives leads to robust models

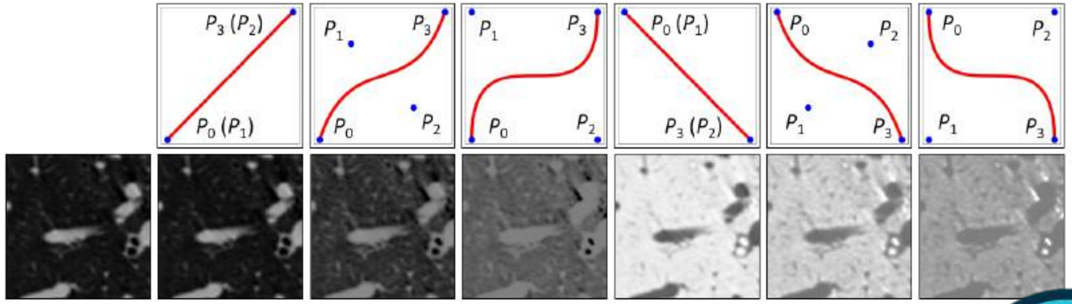

1. Learning organ appearance via **non-linear transformation**

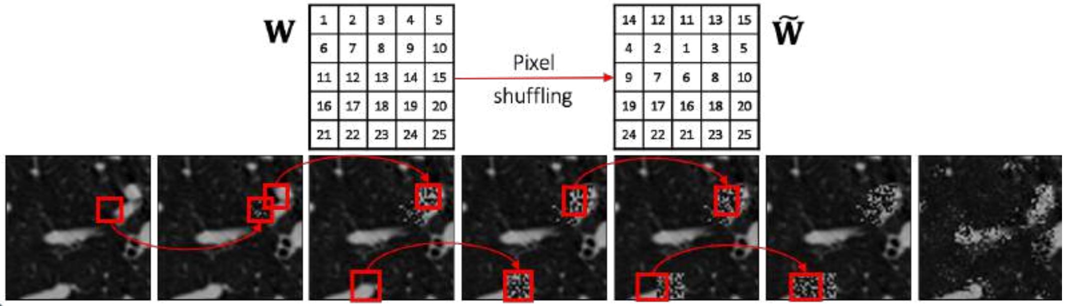

2. Learning organ texture and local boundaries via **local pixel shuffling**

3. Learning organ spatial layout and global geometry via **outer-cutout**

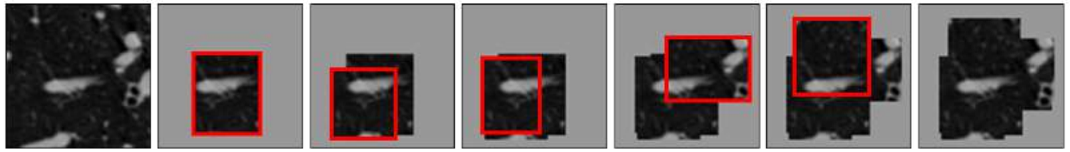

4. Learning local continuities of organs via **inner-cutout**

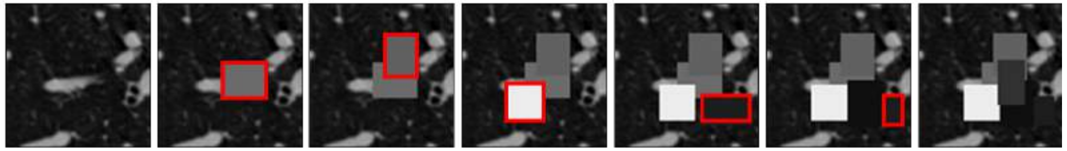

Aim #3: Extract generic knowledge directly from unannotated images

Contribution: Build generic pre-trained 3D models, named “Models Genesis”

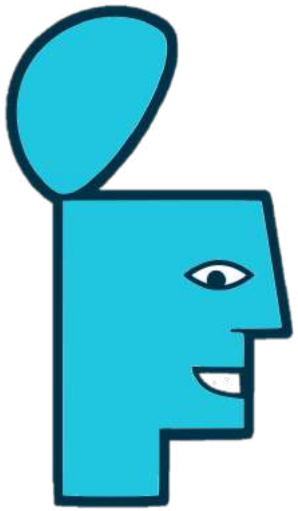

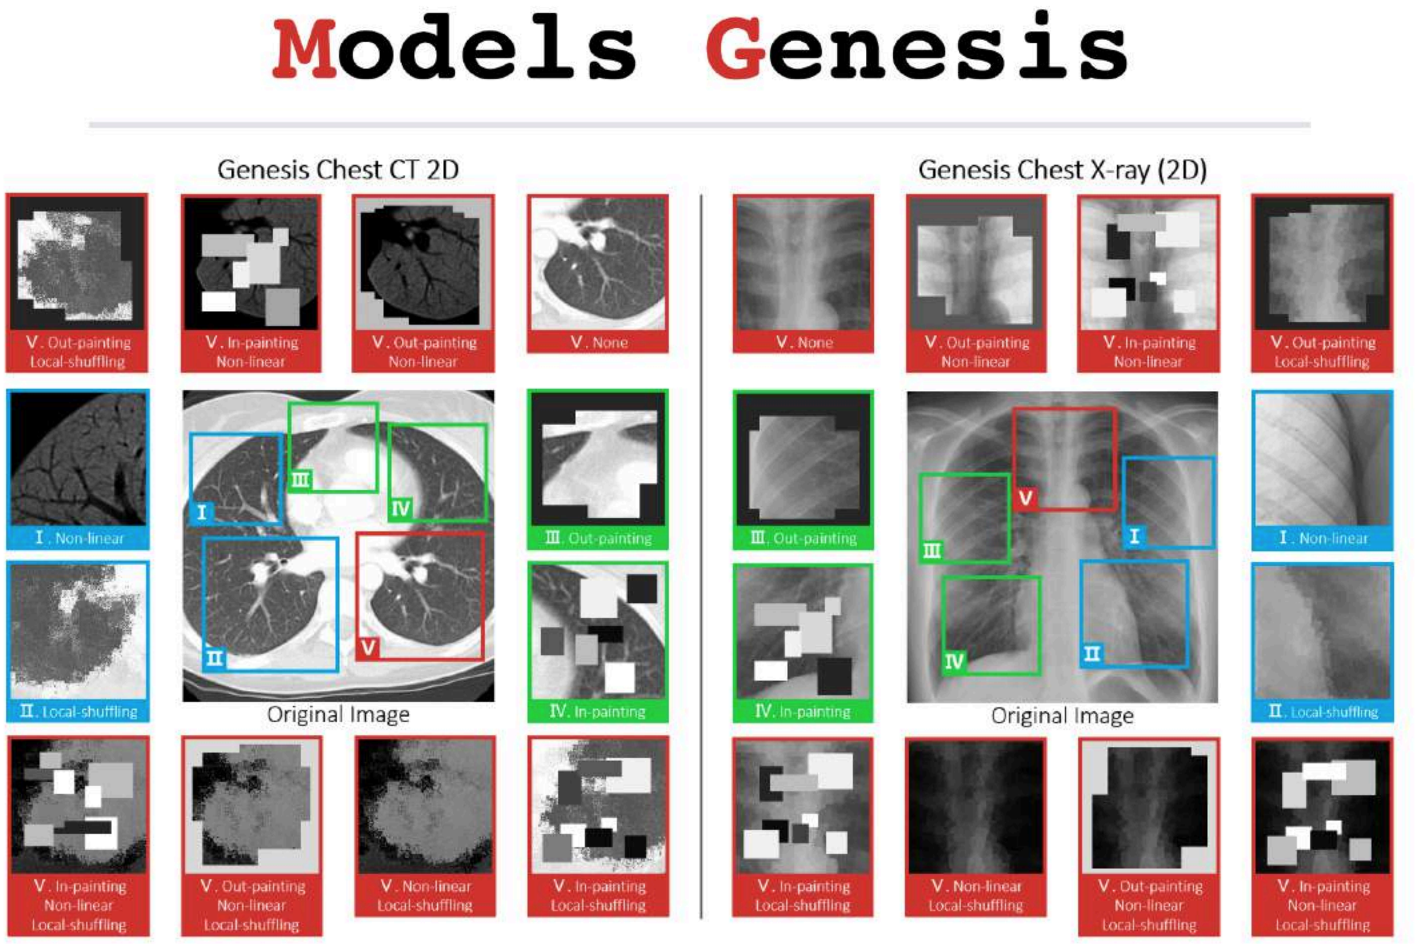

- Z. Zhou, *et al*. Models Genesis. Submitted to *Medical Image Analysis*, 2020. (IF=8.88)
- Z. Zhou, *et al*. Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis. *MICCAI’19.* (Oral; Young Scientist Award)

Aim #3: Extract generic knowledge directly from unannotated images

Contribution: Models Genesis exceed publicly available pre-trained 3D models

Approach	Target tasks				
	NCC[1] (%)	NCS[2] (%)	ECC[3] (%)	LCS[4] (%)	BMS[5] (%)
Random with Uniform Init	94.74±1.97	75.48±0.43	80.36±3.58	78.68±4.23	60.79±1.60
Random with Xavier Init (Glorot and Bengio, 2010)	94.25±5.07	74.05±1.97	79.99±8.06	77.82±3.87	58.52±2.61
Random with MSRA Init (He et al., 2015)	96.03±1.82	76.44±0.45	78.24±3.60	79.76±5.43	63.00±1.73
I3D (Carreira and Zisserman, 2017)	98.26±0.27	71.58±0.55	80.55±1.11	70.65±4.26	67.83±0.75
NiftyNet (Gibson et al., 2018b)	94.14±4.57	52.98±2.05	77.33±8.05	83.23±1.05	60.78±1.60
MedicalNet (Chen et al., 2019b)	95.80±0.49	75.68±0.32	86.43±1.44	**85.52±0.58**[†]	66.09±1.35
De-noising (revised in 3D) (Vincent et al., 2010)	95.92±1.83	73.99±0.62	85.14±3.02	84.36±0.96	57.83±1.57
Patch shuffling (revised in 3D) (Chen et al., 2019a)	91.93±2.32	75.74±0.51	82.15±3.30	82.82±2.35	52.95±6.92
Rubik's Cube (revised) (Zhuang et al., 2019)	96.24±1.27	72.87±0.16	80.49±4.64	75.59±0.20	62.75±1.93
Genesis Chest CT (ours)	**98.34±0.44**	**77.62±0.64**	**87.20±2.87**	85.10±2.15	**67.96±1.29**

[1]NCC Lung nodule false positive reduction in CT images
[2]NCS Lung nodule segmentation in CT images
[3]ECC Pulmonary embolism false positive reduction in CT images
[4]LCS Liver segmentation in CT images
[5]BMS Brain tumor segmentation in MR images

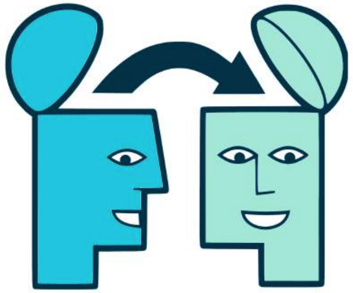

- Z. Zhou, *et al*. Models Genesis. Submitted to *Medical Image Analysis*, 2020. (IF=8.88)
- Z. Zhou, *et al*. Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis. *MICCAI'19*. (Oral; Young Scientist Award)

Aim #3: Extract generic knowledge directly from unannotated images

Contribution: Models Genesis consistently top any 2D approaches

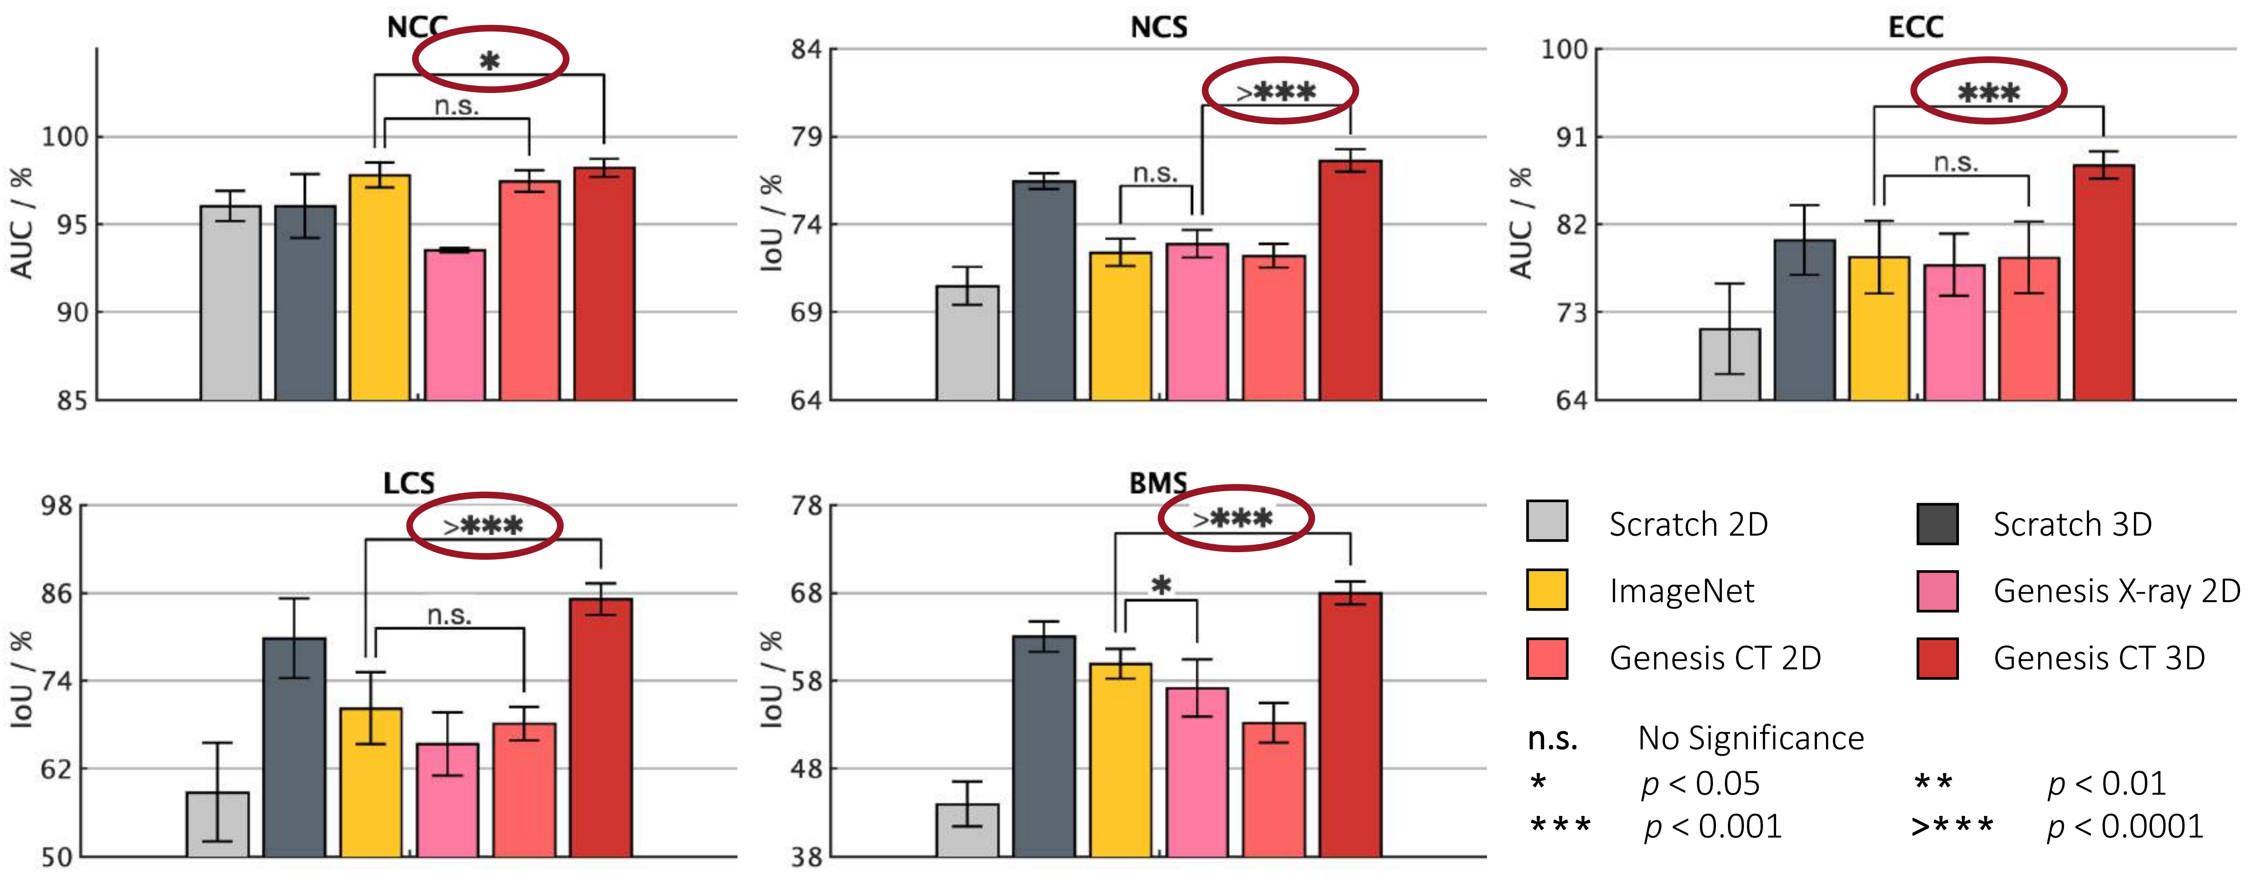

- Z. Zhou, *et al*. Models Genesis. Submitted to *Medical Image Analysis*, 2020. (IF=8.88)
- Z. Zhou, *et al*. Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis. *MICCAI'19.* (Oral; Young Scientist Award)

Aim #3: Extract generic knowledge directly from unannotated images

Proposal: Extend to modality-oriented and organ-oriented models

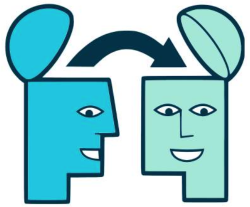

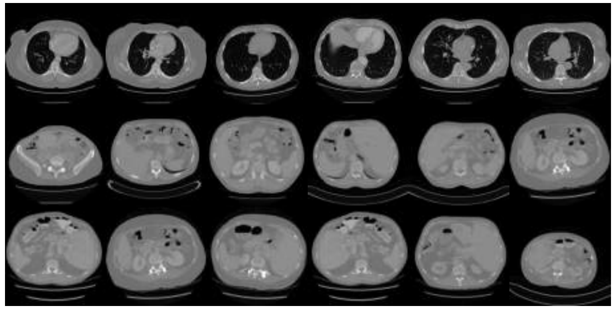

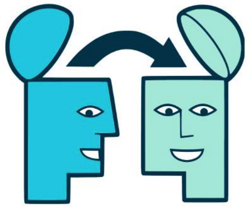

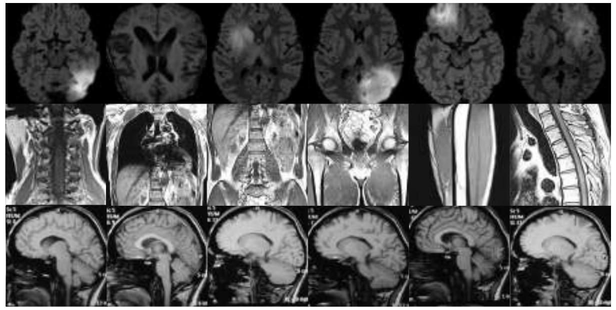

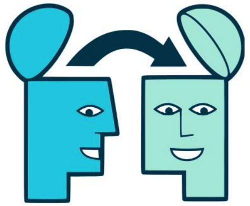

Aim #3: Extract generic knowledge directly from unannotated images

Proposal: Extend to modality-oriented and organ-oriented models

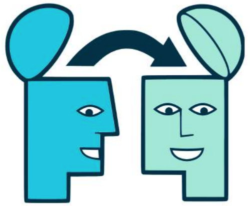

Genesis X-ray

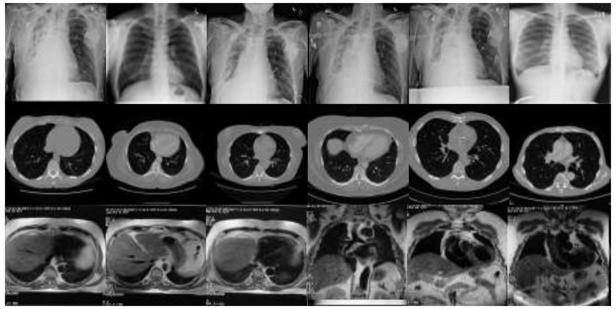

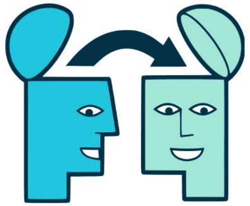

Genesis Lung

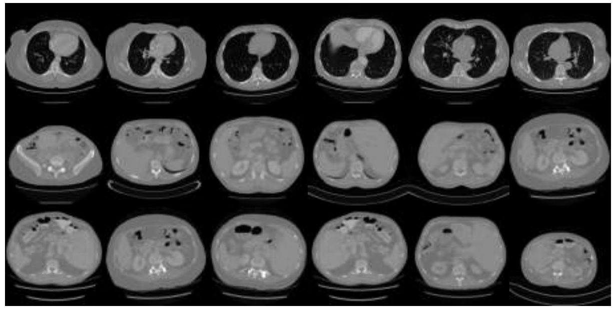

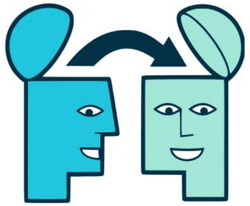

Genesis CT

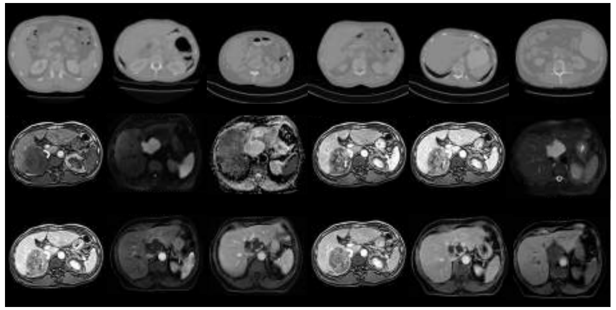

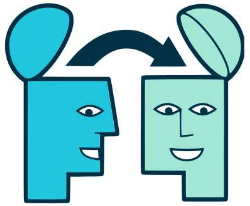

Genesis Liver

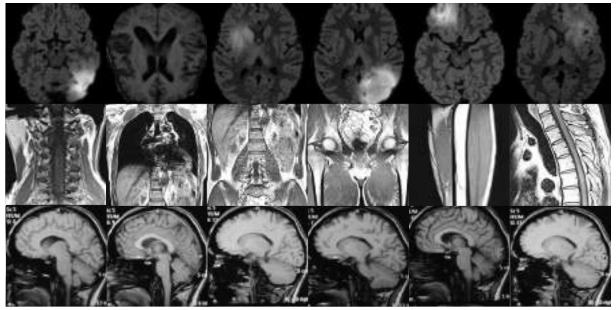

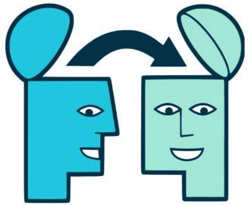

Genesis MRI

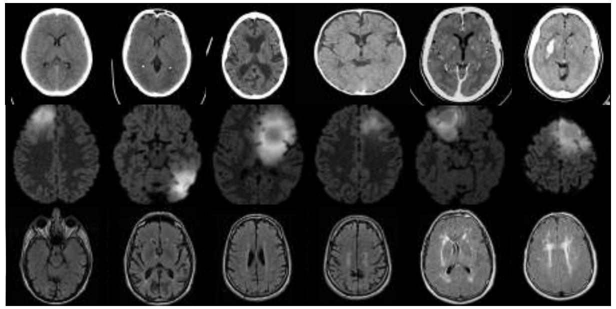

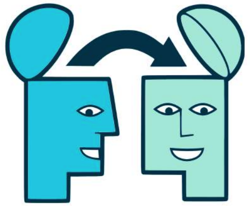

Genesis Brain

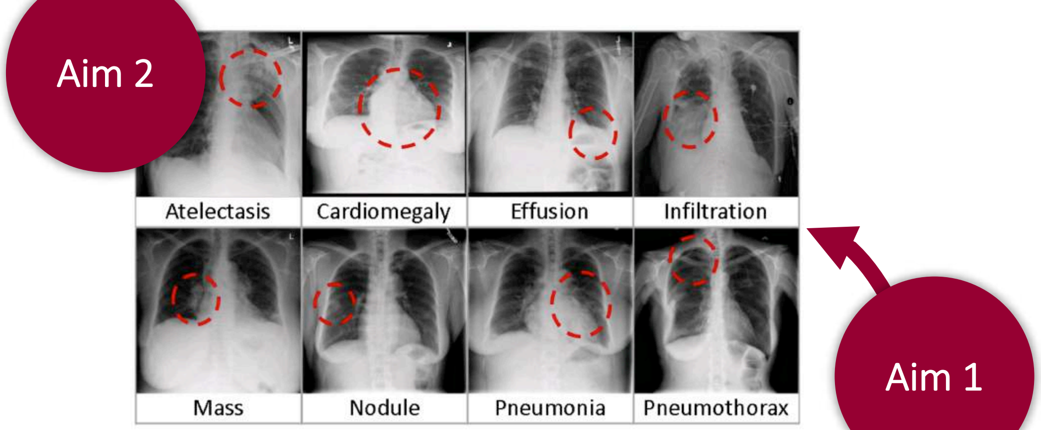

$ 1,000 annotation budget ☺

$ 1,000,000 annotation cost ☹

Research goal: Exploit novel methods to minimize the manual labeling efforts for a rapid, precise computer-aided diagnosis of lung diseases

Aim #1: Acquire necessary annotation efficiently from human experts

Aim #2: Utilize existing annotation effectively from advanced architecture

Aim #3: Extract generic knowledge directly from unannotated images

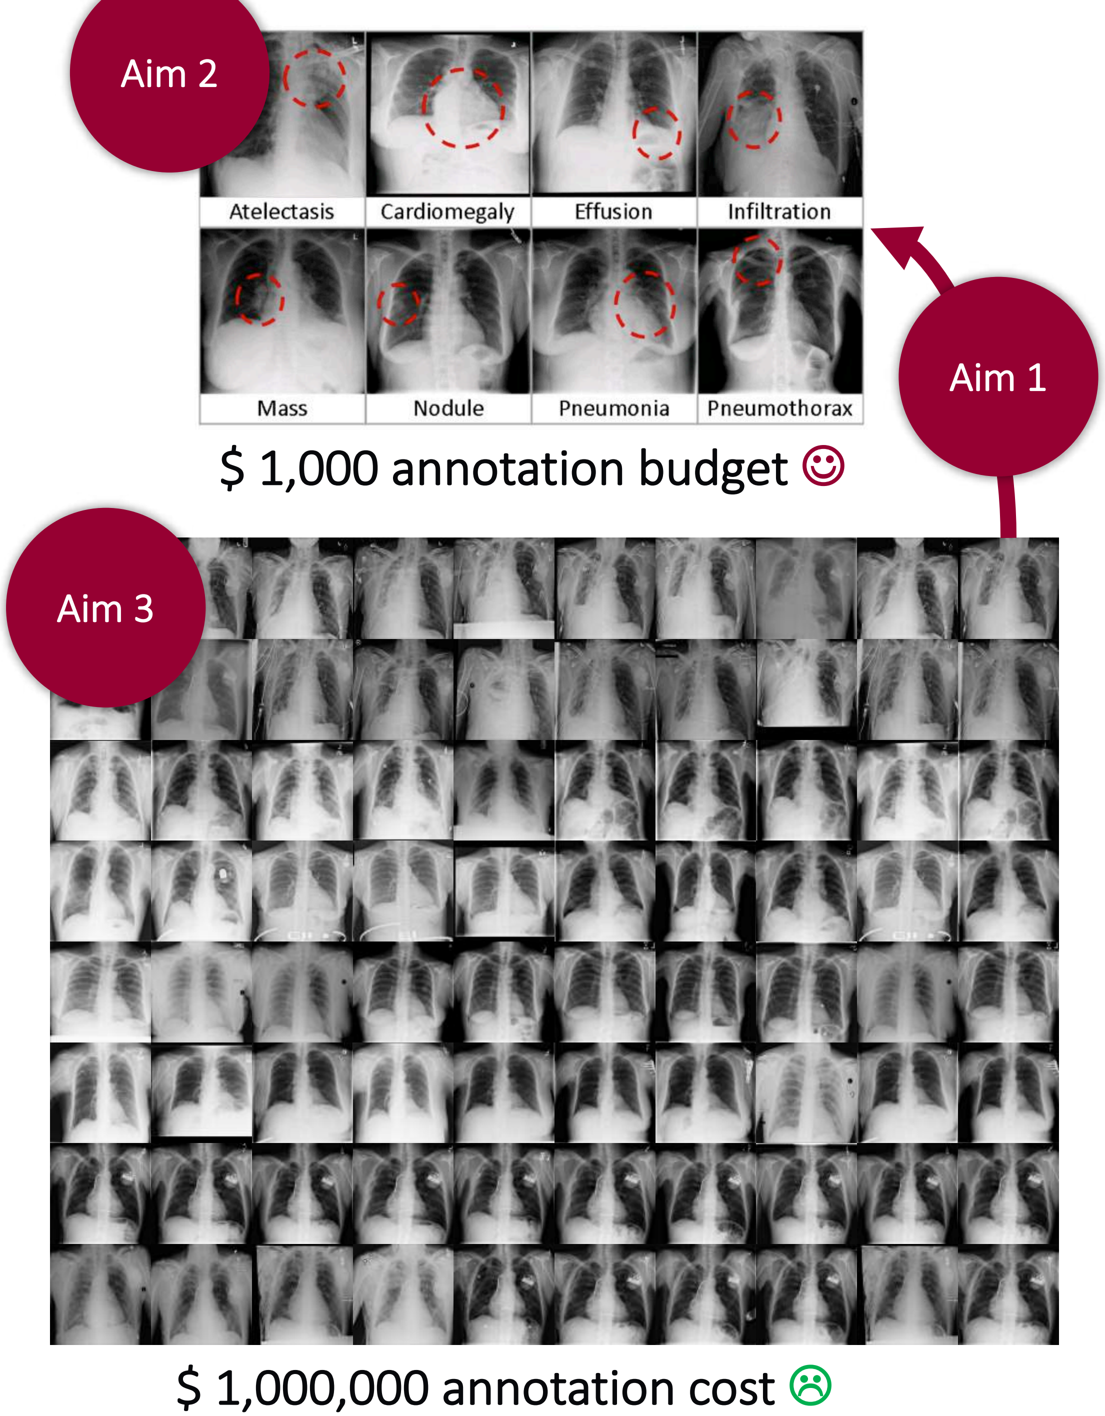

Research goal: Exploit novel methods to minimize the manual labeling efforts for a rapid, precise computer-aided diagnosis of lung diseases

Aim #1: Acquire necessary annotation efficiently from human experts

Aim #2: Utilize existing annotation effectively from advanced architecture

Aim #3: Extract generic knowledge directly from unannotated images

Clinical application: COVID-19

Dataset: Consist of more than 200 positive patients and thousands of negative patients, provided by Renmin Hospital of Wuhan University

Goal: Create a reliable CAD system in a short span of time for automated COVID-19 diagnosis

Acknowledgements

- Jianming Liang, Ph.D.
- Edward H. Shortliffe, M.D., Ph.D.
- Murthy Devarakonda, Ph.D.
- Michael B. Gotway, M.D.

Cost-Effective Deep Learning in Medical Image Analysis

Zongwei Zhou

College of Health Solutions, Arizona State University

P: 1-(480)738-2575 | E: zongweiz@asu.edu