Annotation-efficient Deep Learning for Computer-aided Diagnosis in Medical Imaging

Zongwei Zhou

Ph.D. Candidate, Biomedical Informatics

College of Health Solutions, Arizona State University

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

Introduction

Objective

Aim #1

Aim #2

Aim #3

Summary

Deep Learning propels us into the so-called artificial intelligence (AI) era

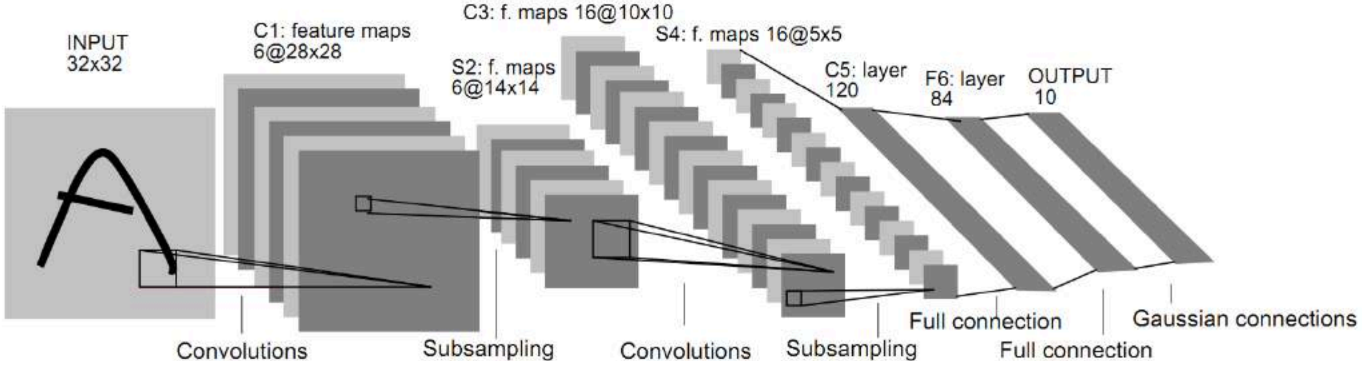

Imaging data account for about 90% of all healthcare data

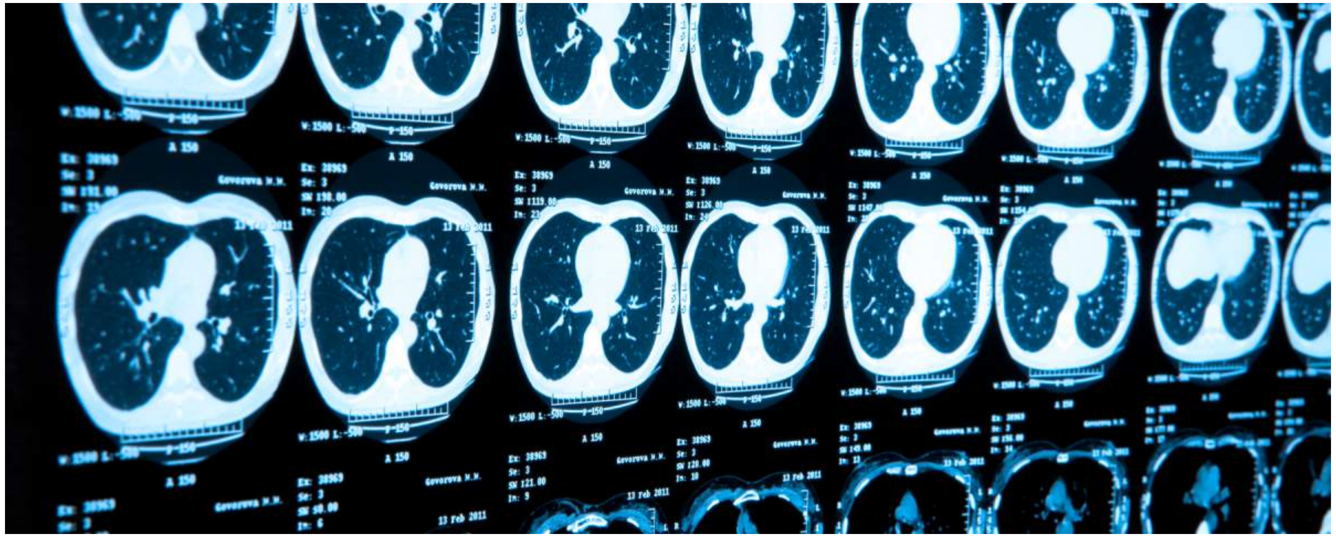

1. LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. "Deep learning." nature 521.7553 (2015): 436-444
2. "The Digital Universe Driving Data Growth in Healthcare." published by EMC with research and analysis from IDC (12/13)

Introduction

Objective

Aim #1

Aim #2

Aim #3

Summary

Deep Learning works well in medical imaging, but it demands massive annotation costs.

To match human diagnostic precision, deep learning algorithms require

- **42,290** radiologist-labeled CT images for lung cancer diagnosis
- **128,175** ophthalmologist-labeled retinal images for diabetic retinopathy detection
- **129,450** dermatologist-labeled images for skin cancer classification

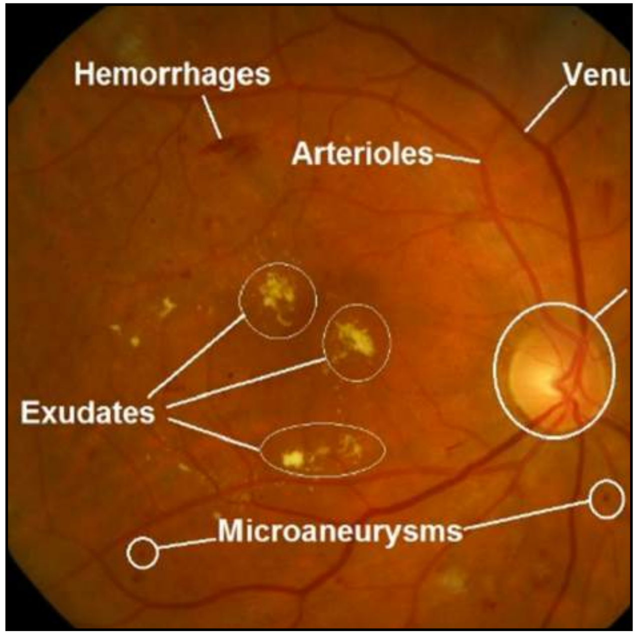

1. Ardila, Diego, et al. "End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography." Nature medicine 25.6 (2019): 954-961.
2. Gulshan, Varun, et al. "Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs." Jama 316.22 (2016): 2402-2410.
3. Esteva, Andre, et al. "Dermatologist-level classification of skin cancer with deep neural networks." nature 542.7639 (2017): 115-118.

Deep Learning works well in medical imaging, but it demands massive annotation costs.

To match human diagnostic precision, deep learning algorithms require

- **42,290** radiologist-labeled CT images for lung cancer diagnosis
- **128,175** ophthalmologist-labeled retinal images for diabetic retinopathy detection
- **129,450** dermatologist-labeled images for skin cancer classification

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

Consider the scenarios as follows:

- 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

Computer-aided diagnosis of rare diseases or rapid response to global pandemics are severely under-explored owing to the difficulty of collecting a sizeable amount labeled data.

Introduction

Objective

Aim #1

Aim #2

Aim #3

Summary

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

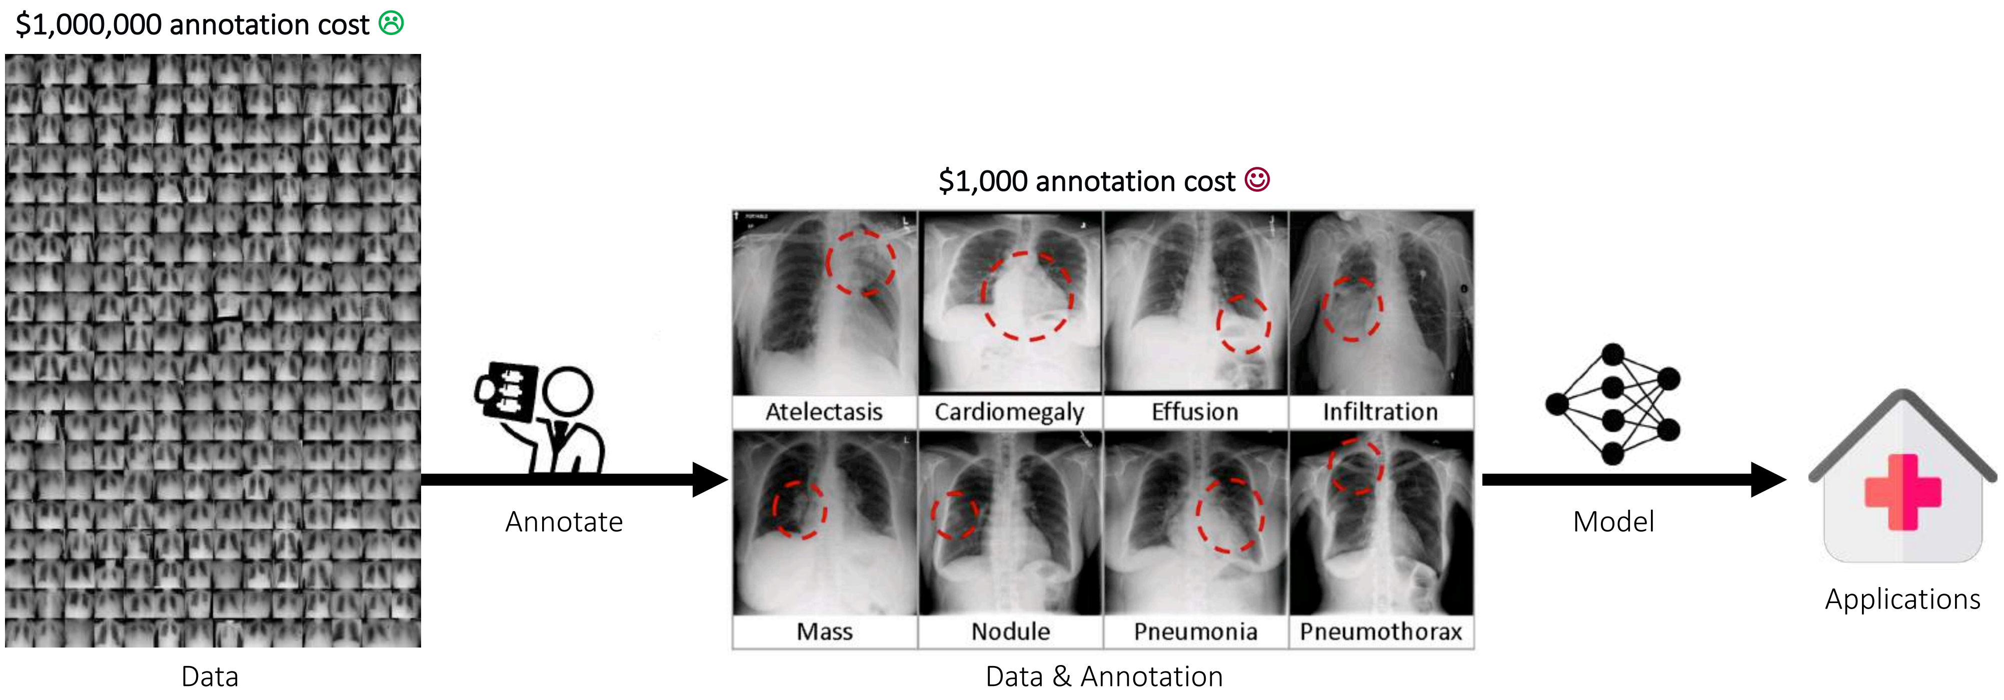

Research goal: Exploit novel methods to minimize the manual labeling efforts for a rapid, precise computer-aided diagnosis system

Aim #1: Acquiring necessary annotation efficiently from human experts

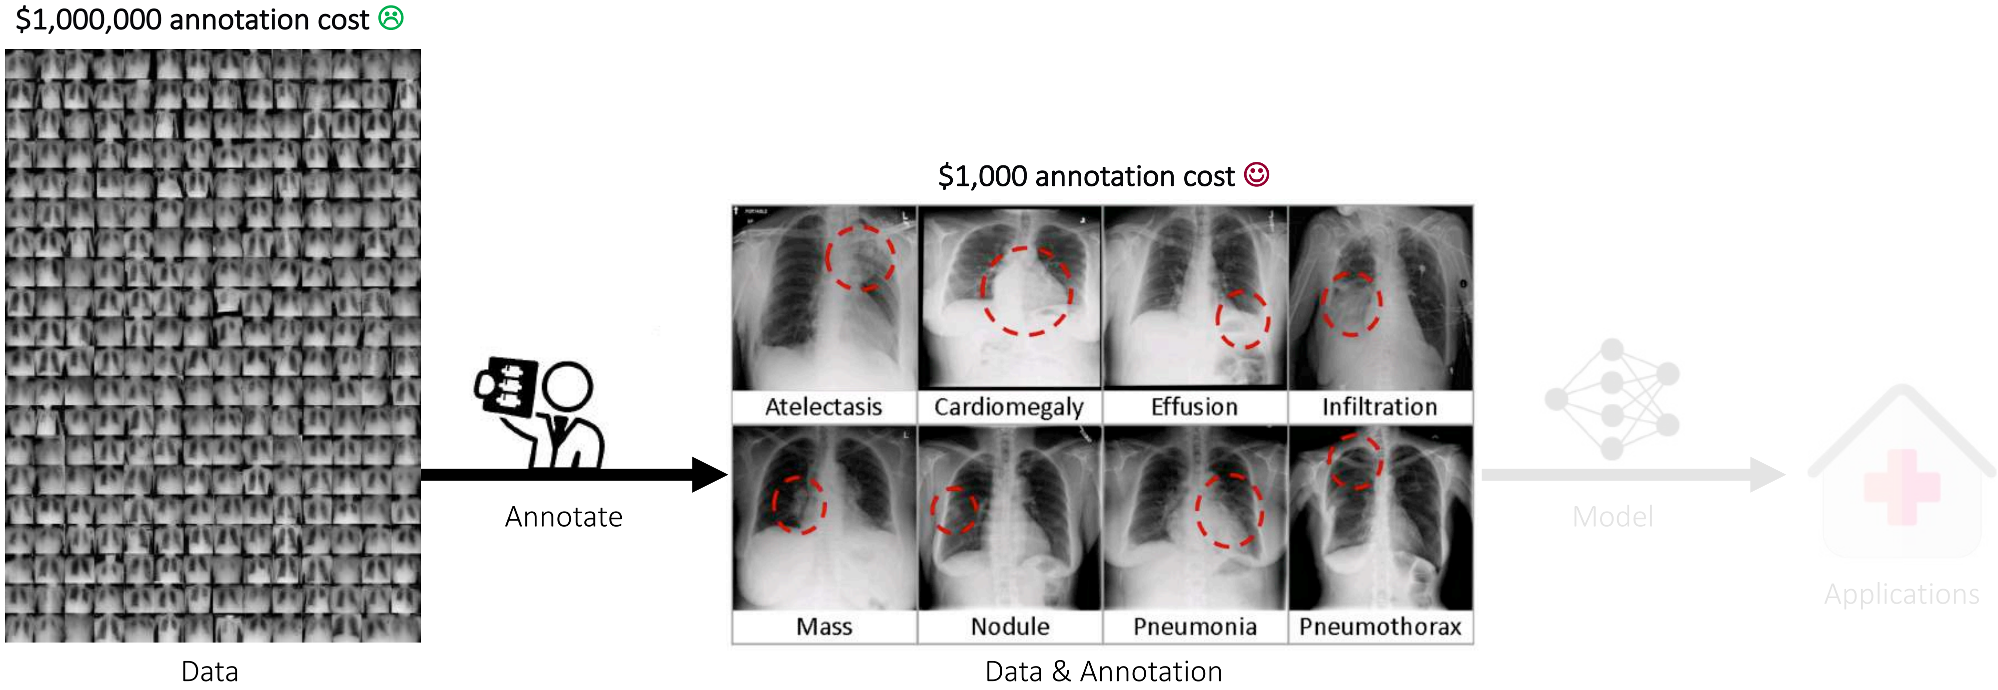

Research goal: Exploit novel methods to minimize the manual labeling efforts for a rapid, precise computer-aided diagnosis system

Aim #1: Acquiring necessary annotation efficiently from human experts

Aim #2: Utilizing existing annotation effectively from advanced architecture

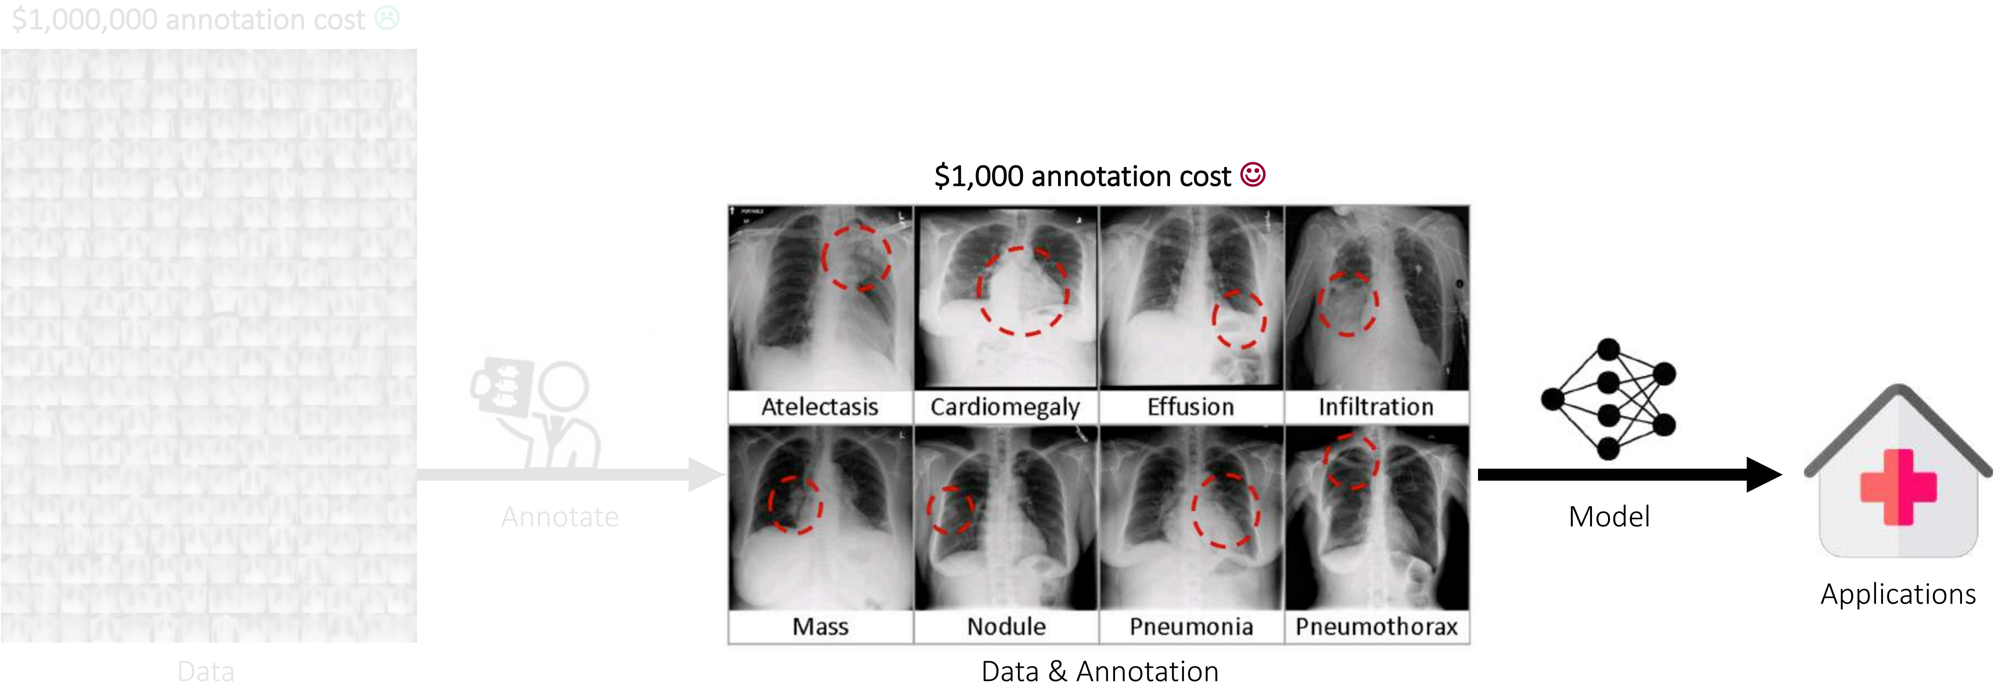

Research goal: Exploit novel methods to minimize the manual labeling efforts for a rapid, precise computer-aided diagnosis system

Aim #1: Acquiring necessary annotation efficiently from human experts

Aim #2: Utilizing existing annotation effectively from advanced architecture

Aim #3: Extracting generic knowledge directly from unannotated images

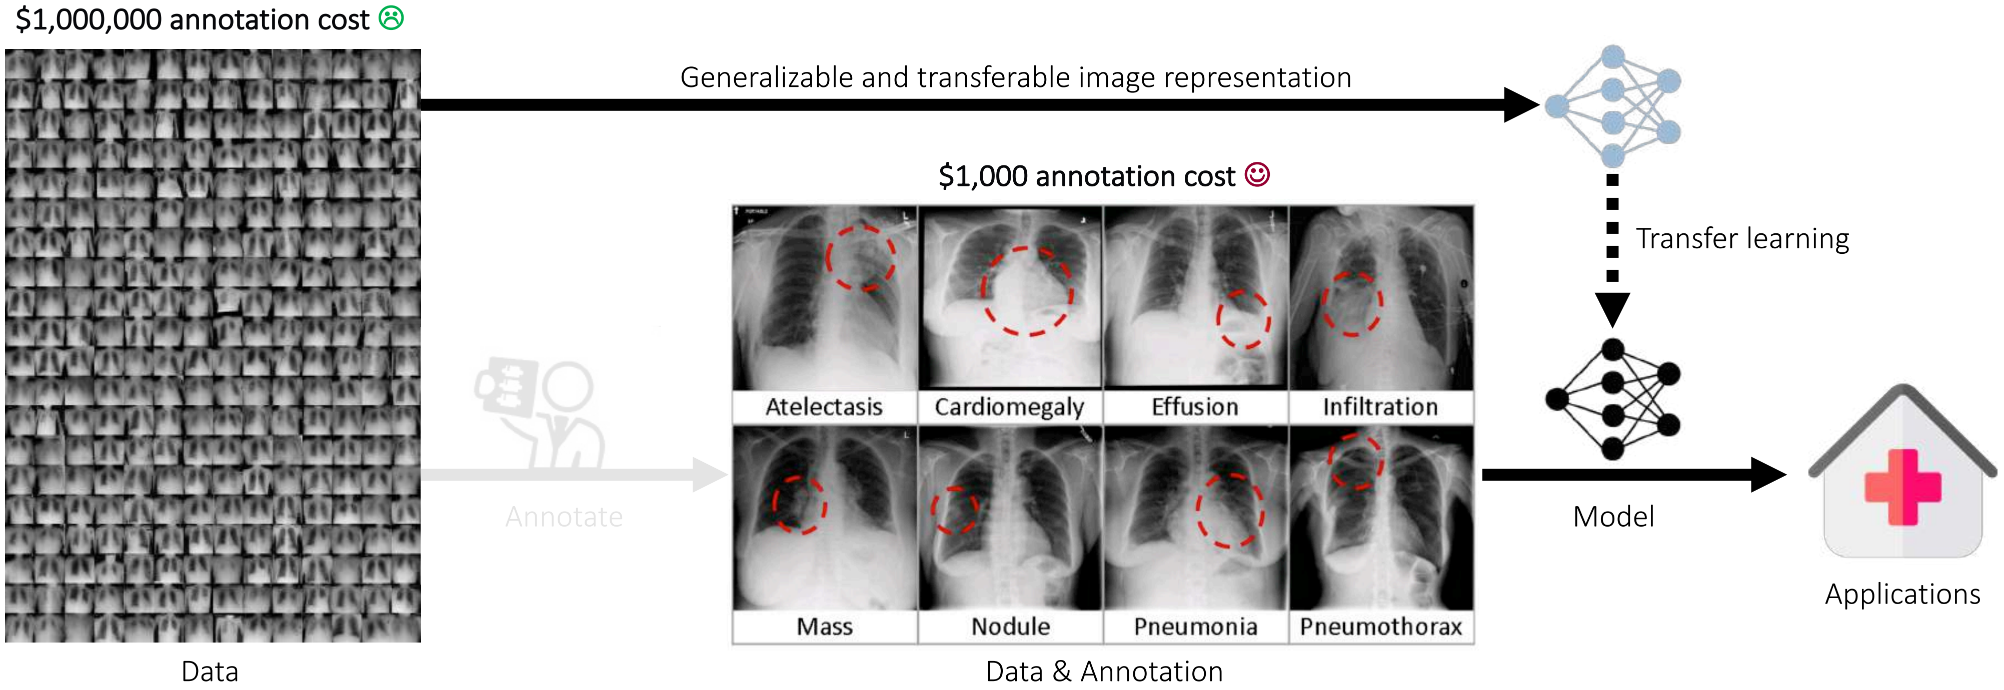

Research goal: Exploit novel methods to minimize the manual labeling efforts for a rapid, precise computer-aided diagnosis system

Aim #1: Acquiring necessary annotation efficiently from human experts

Aim #2: Utilizing existing annotation effectively from advanced architecture

Aim #3: Extracting generic knowledge directly from unannotated images

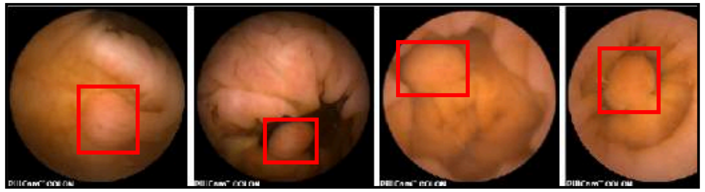
Polyp detection

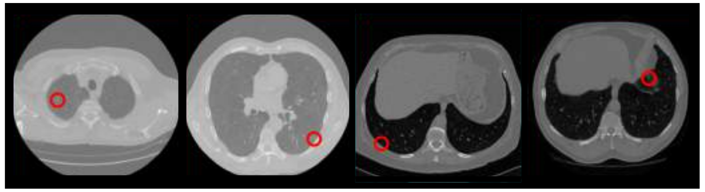
Lung nodule detection

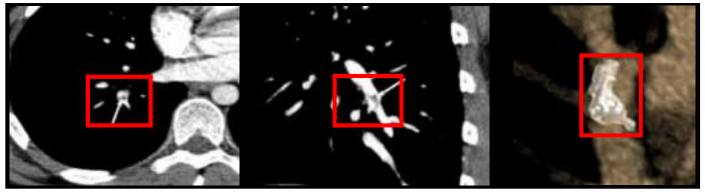
Pulmonary embolism detection

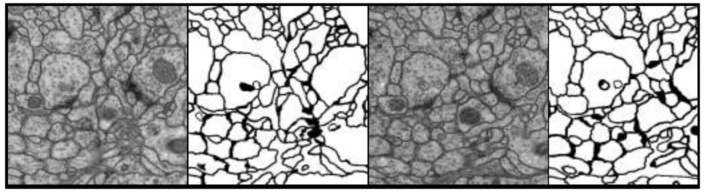
Neuronal structure segmentation

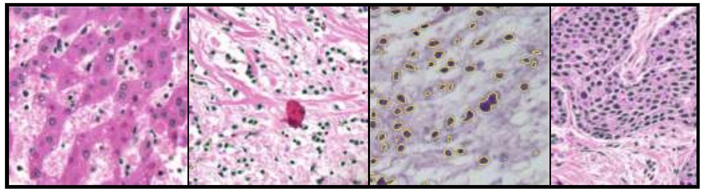
Cell/nuclei segmentation

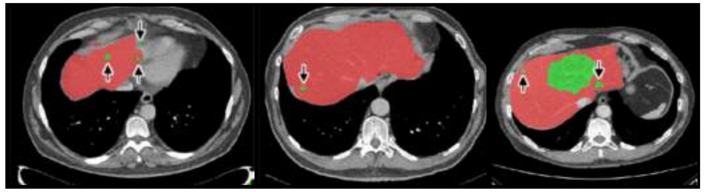
Liver/lesion segmentation

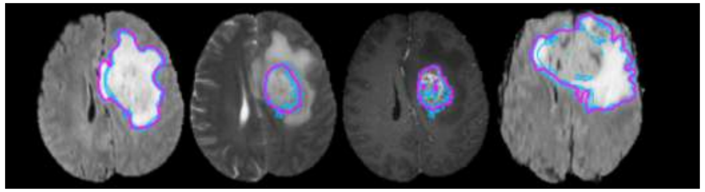
Brain/tumor segmentation


Kidney/lesion segmentation

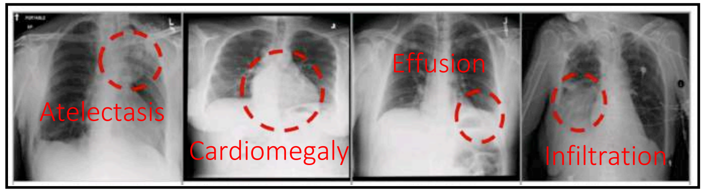

Pulmonary diseases classification

Aim #1: Acquiring necessary annotation efficiently from human experts

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

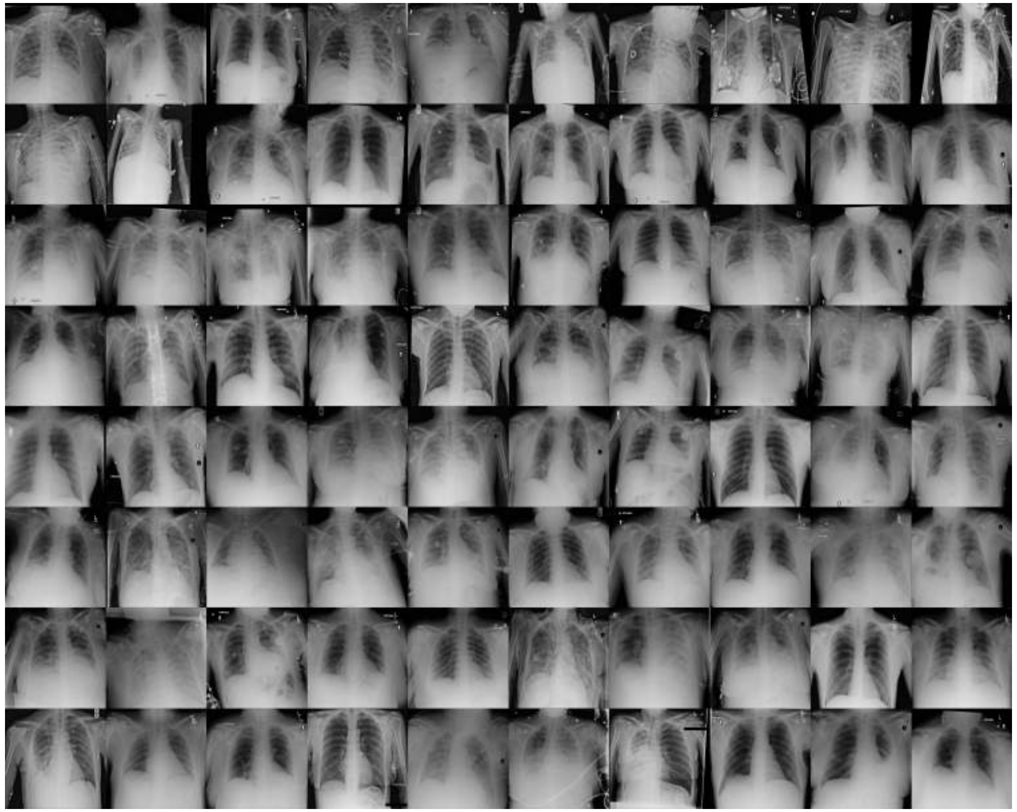

$ 1,000,000 annotation cost ☹

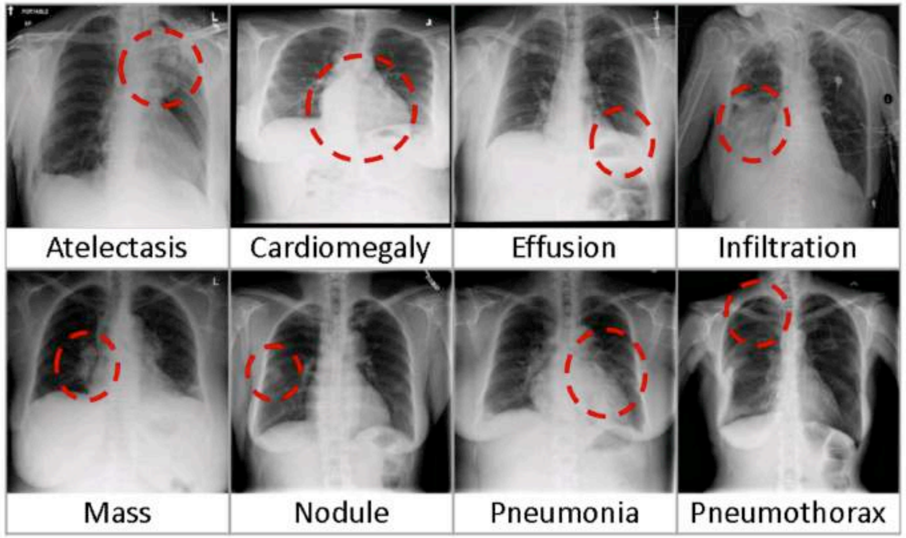

$ 1,000 annotation budget ☺

Aim #1: Acquiring necessary annotation efficiently from human experts

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

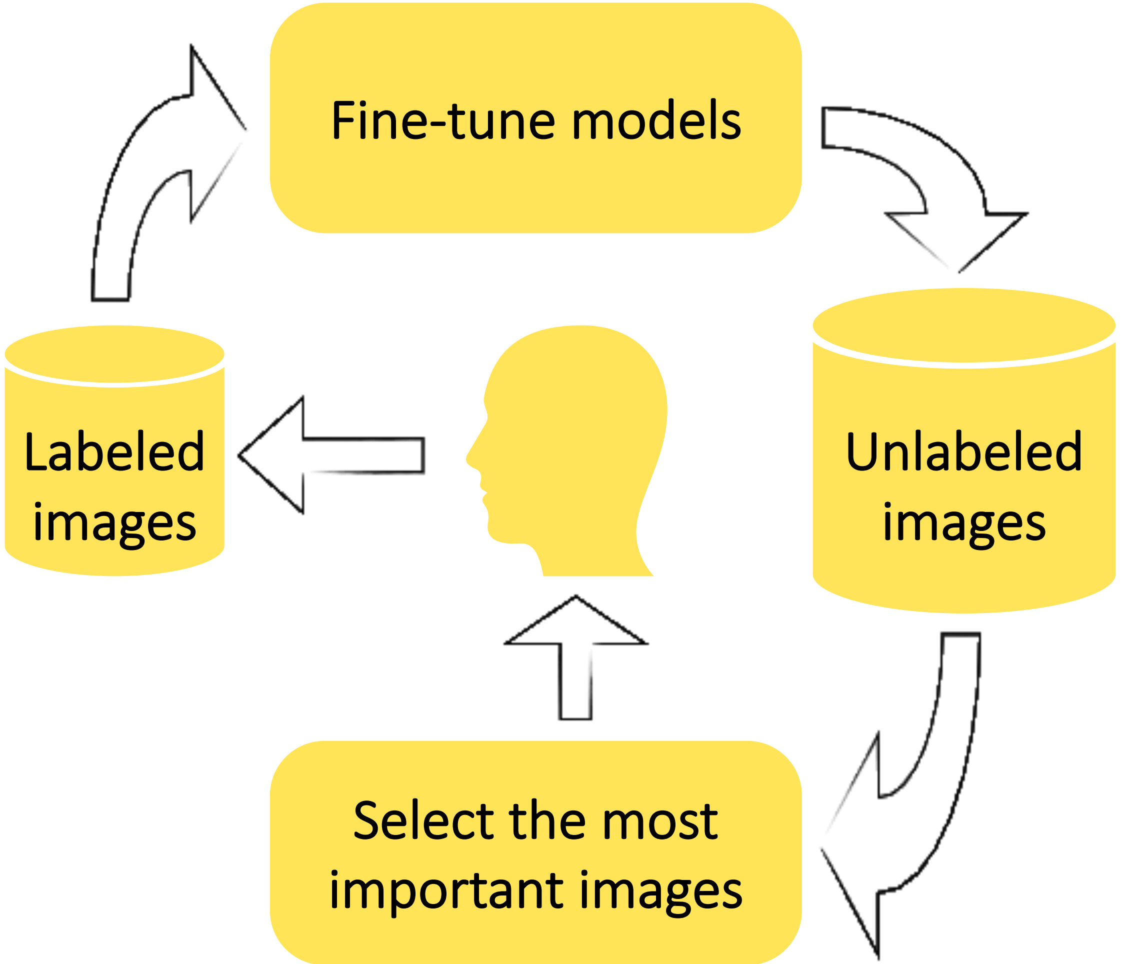

Aim #1: Acquiring necessary annotation efficiently from human experts

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

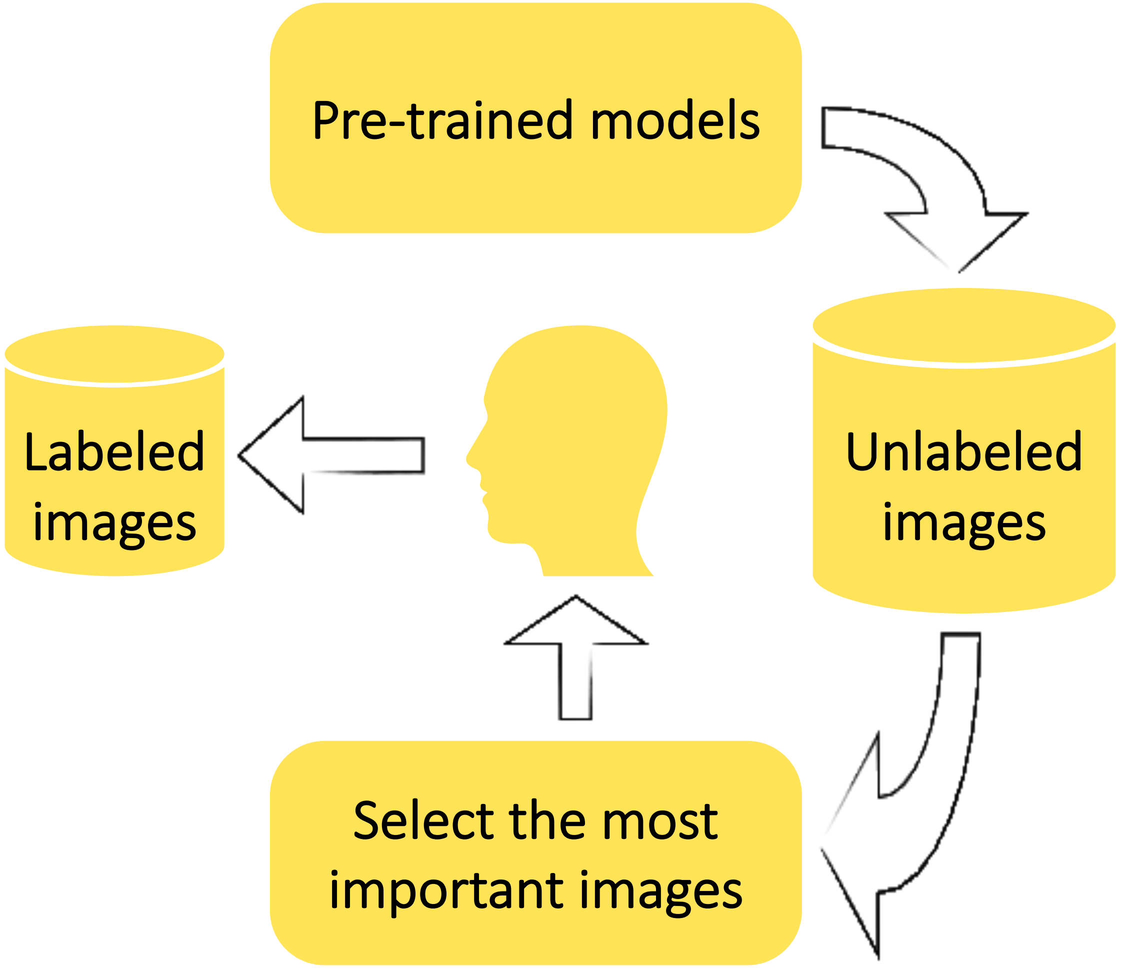

Aim #1: Acquiring necessary annotation efficiently from human experts

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

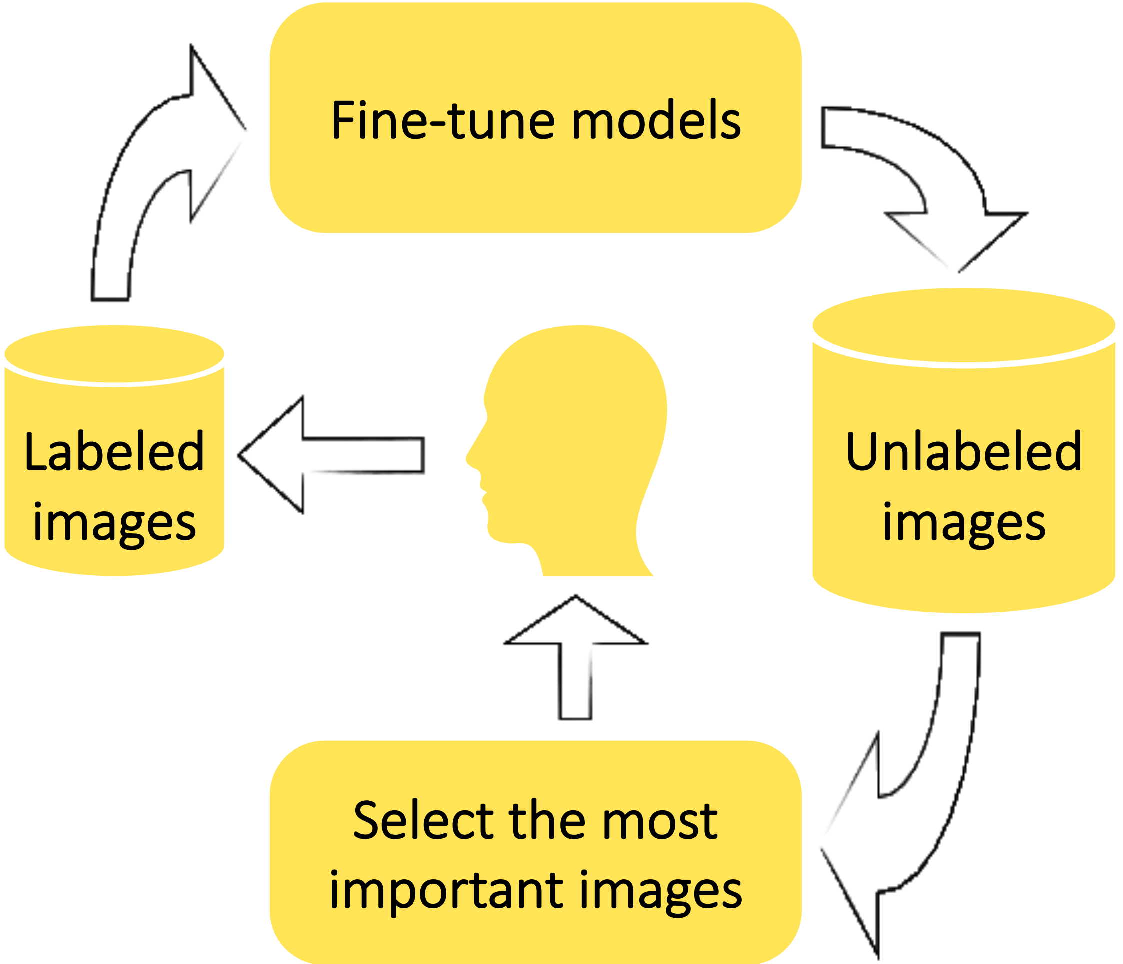

Aim #1: Acquiring necessary annotation efficiently from human experts

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

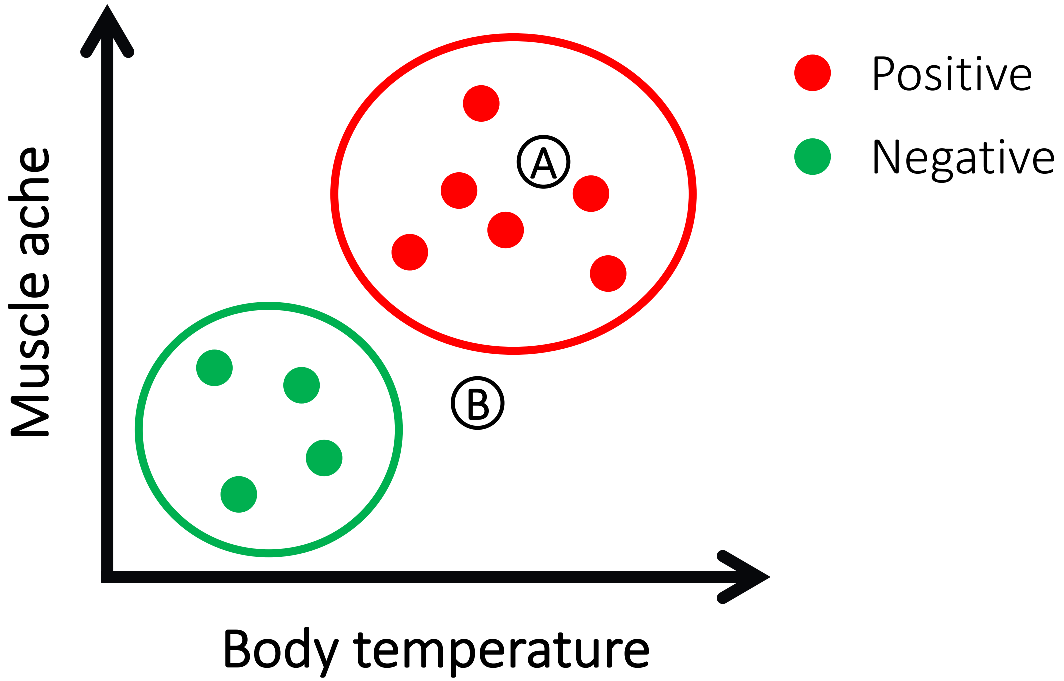

Select the most important samples

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

Aim #1: Acquiring necessary annotation efficiently from human experts

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

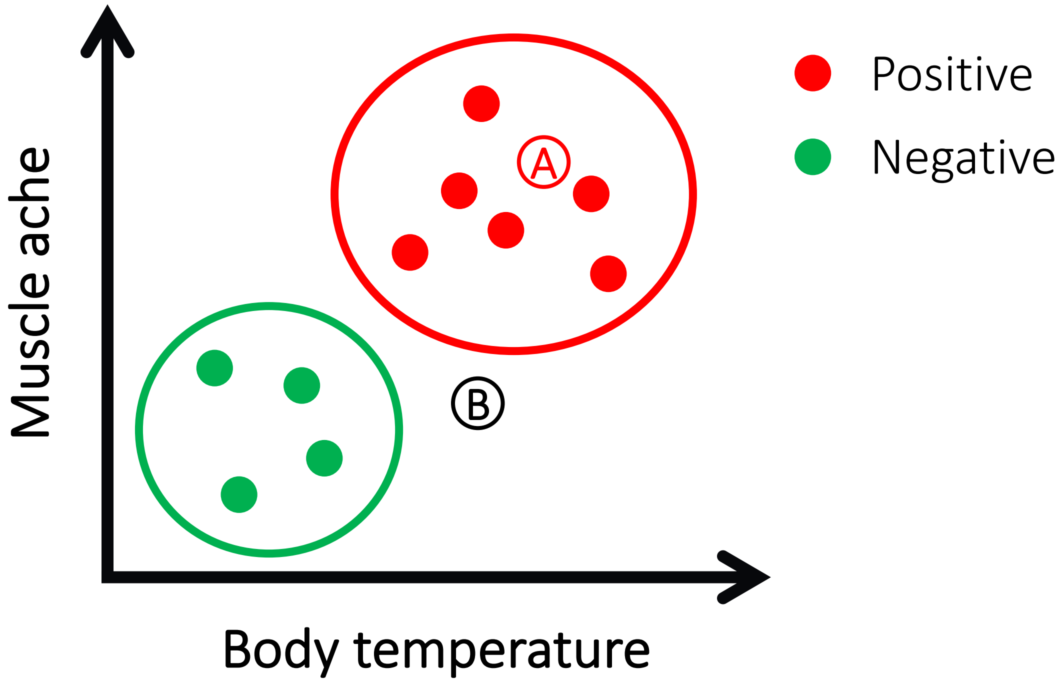

Select the most important samples

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

Aim #1: Acquiring necessary annotation efficiently from human experts

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

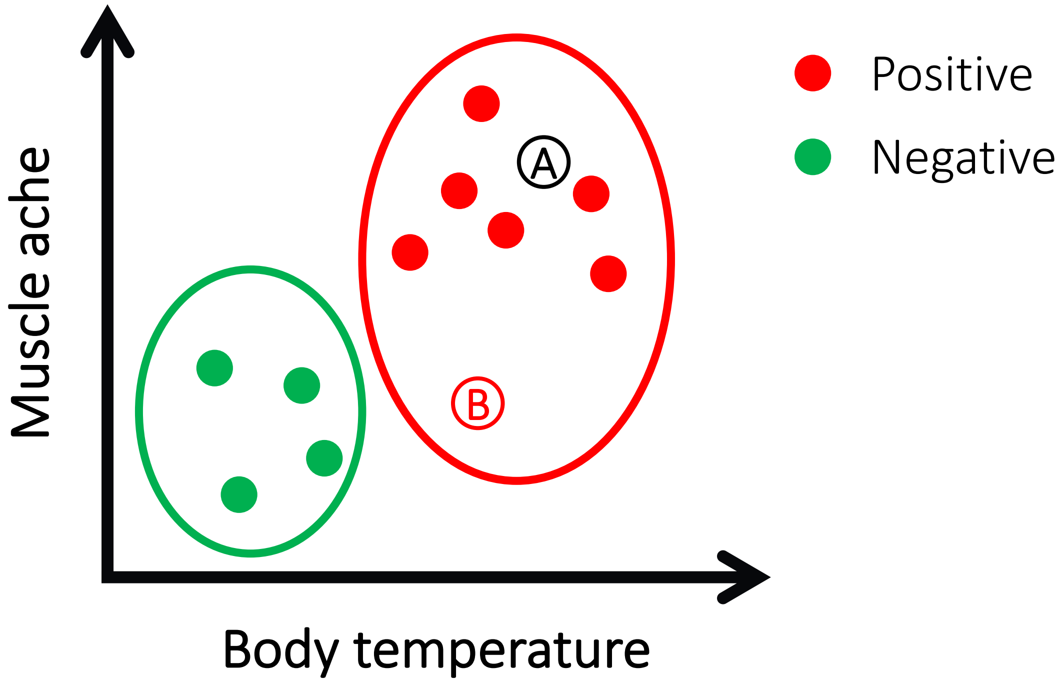

Select the most important samples

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

Introduction

Objective

Aim #1

Aim #2

Aim #3

Summary

Aim #1: Acquiring necessary annotation efficiently from human experts

Approach: “Human-in-the-loop” active learning procedure

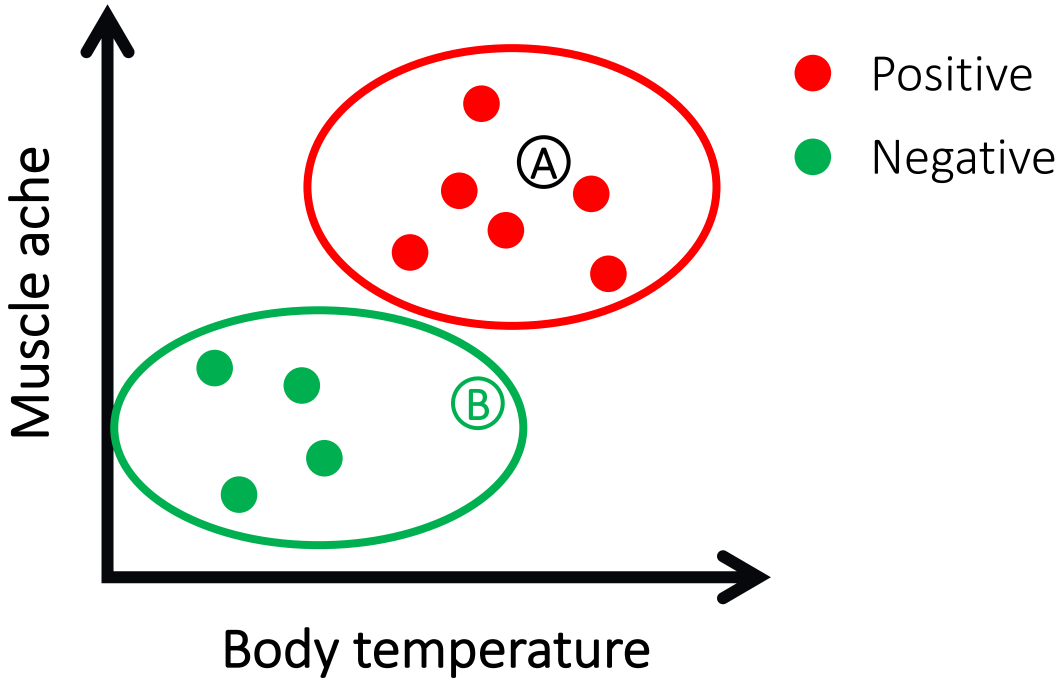

Select the most important samples

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

Aim #1: Acquiring necessary annotation efficiently from human experts

Approach: Active, Continual Fine-Tuning

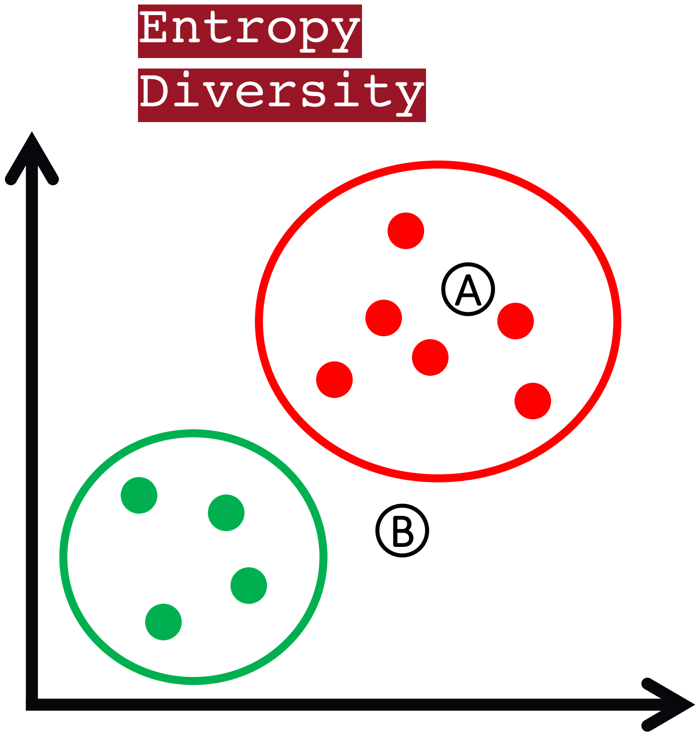

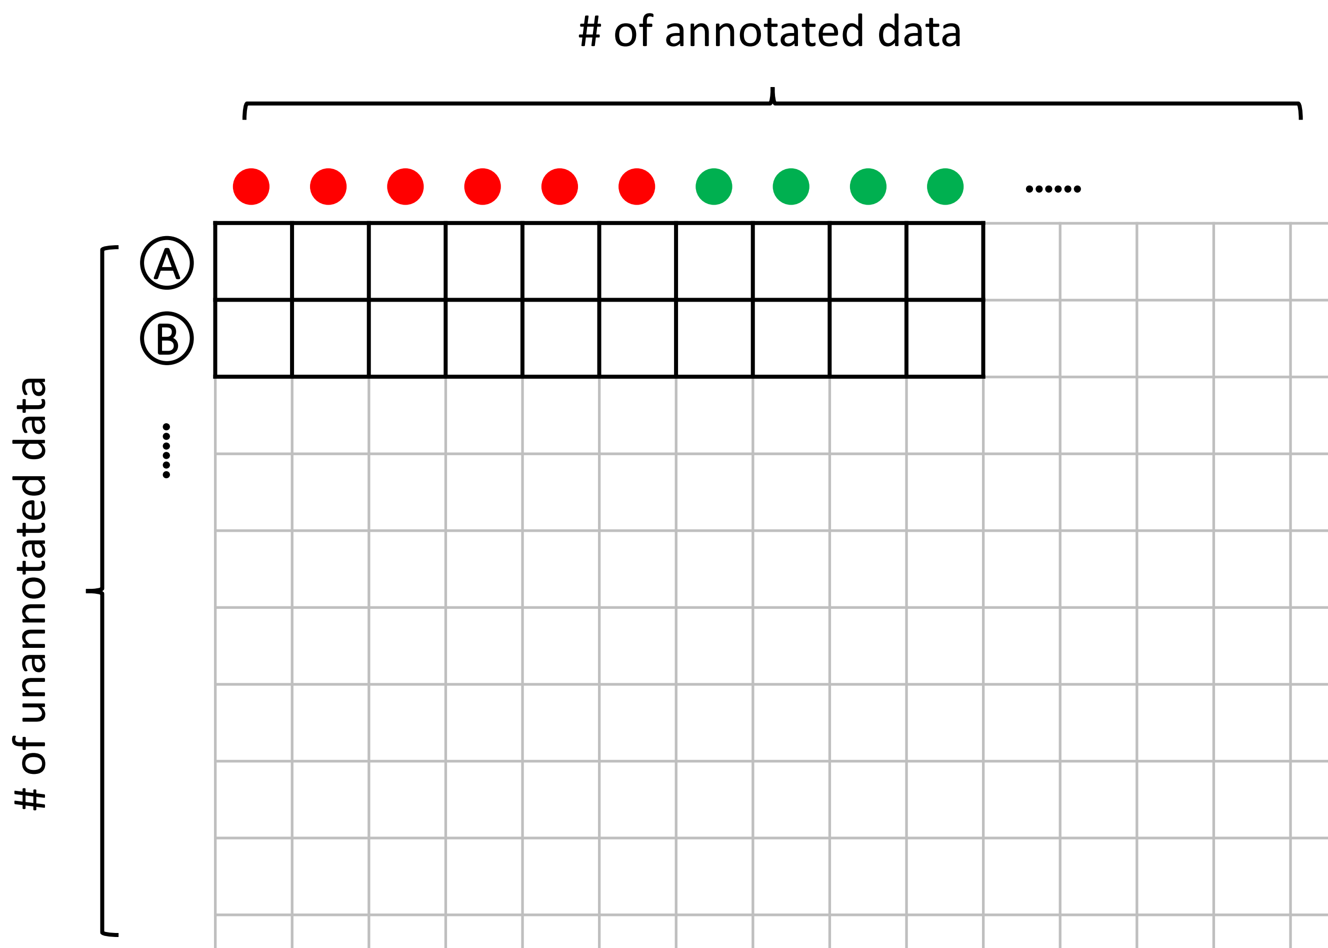

Aim #1: Acquiring necessary annotation efficiently from human experts

Approach: Active, Continual Fine-Tuning

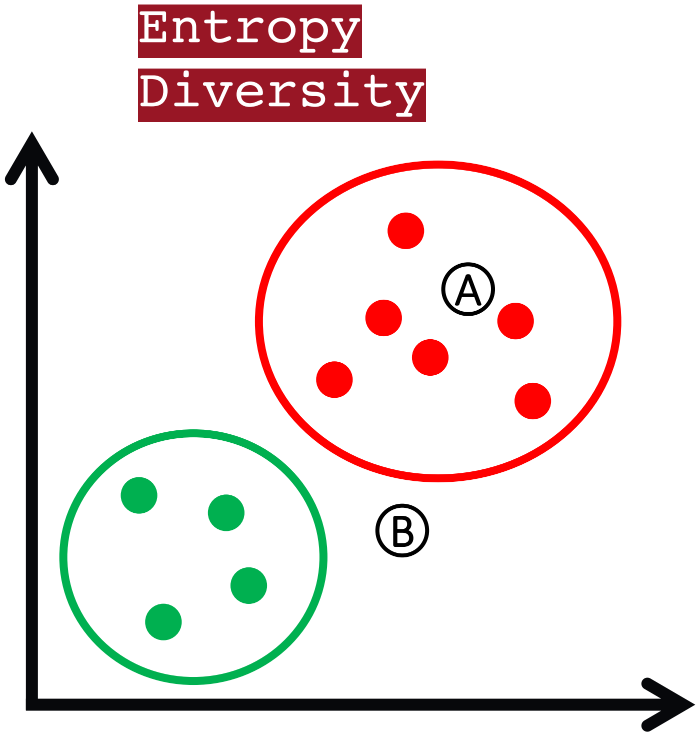

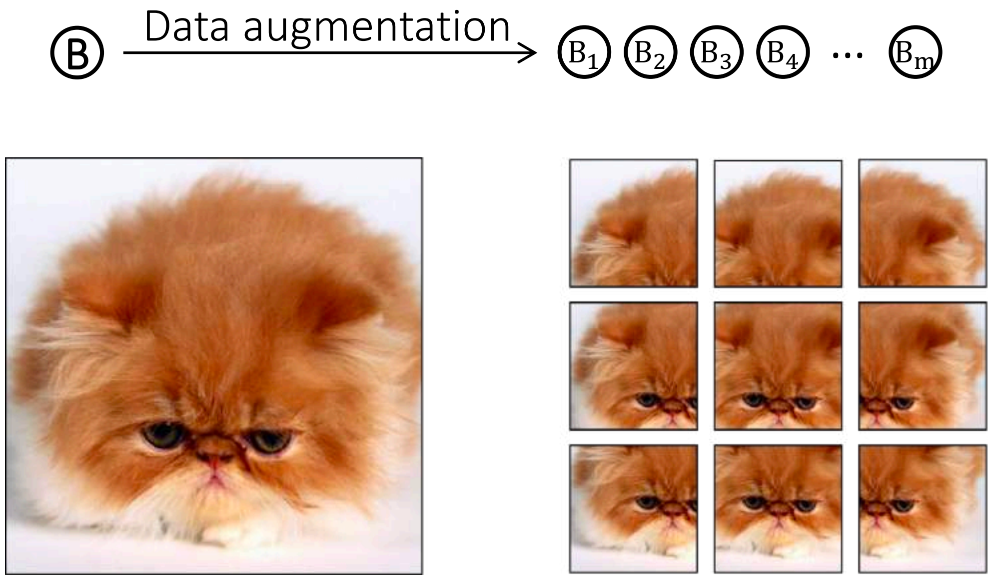

To boost the performance of CNNs, multiple patches are usually generated for each image via **data augmentation**; these patches generated from the same image share the **same label**, and are naturally expected to have **similar predictions** by the current CNN.

Aim #1: Acquiring necessary annotation efficiently from human experts

Hypothesis: Wisely selecting important samples can reduce annotation cost

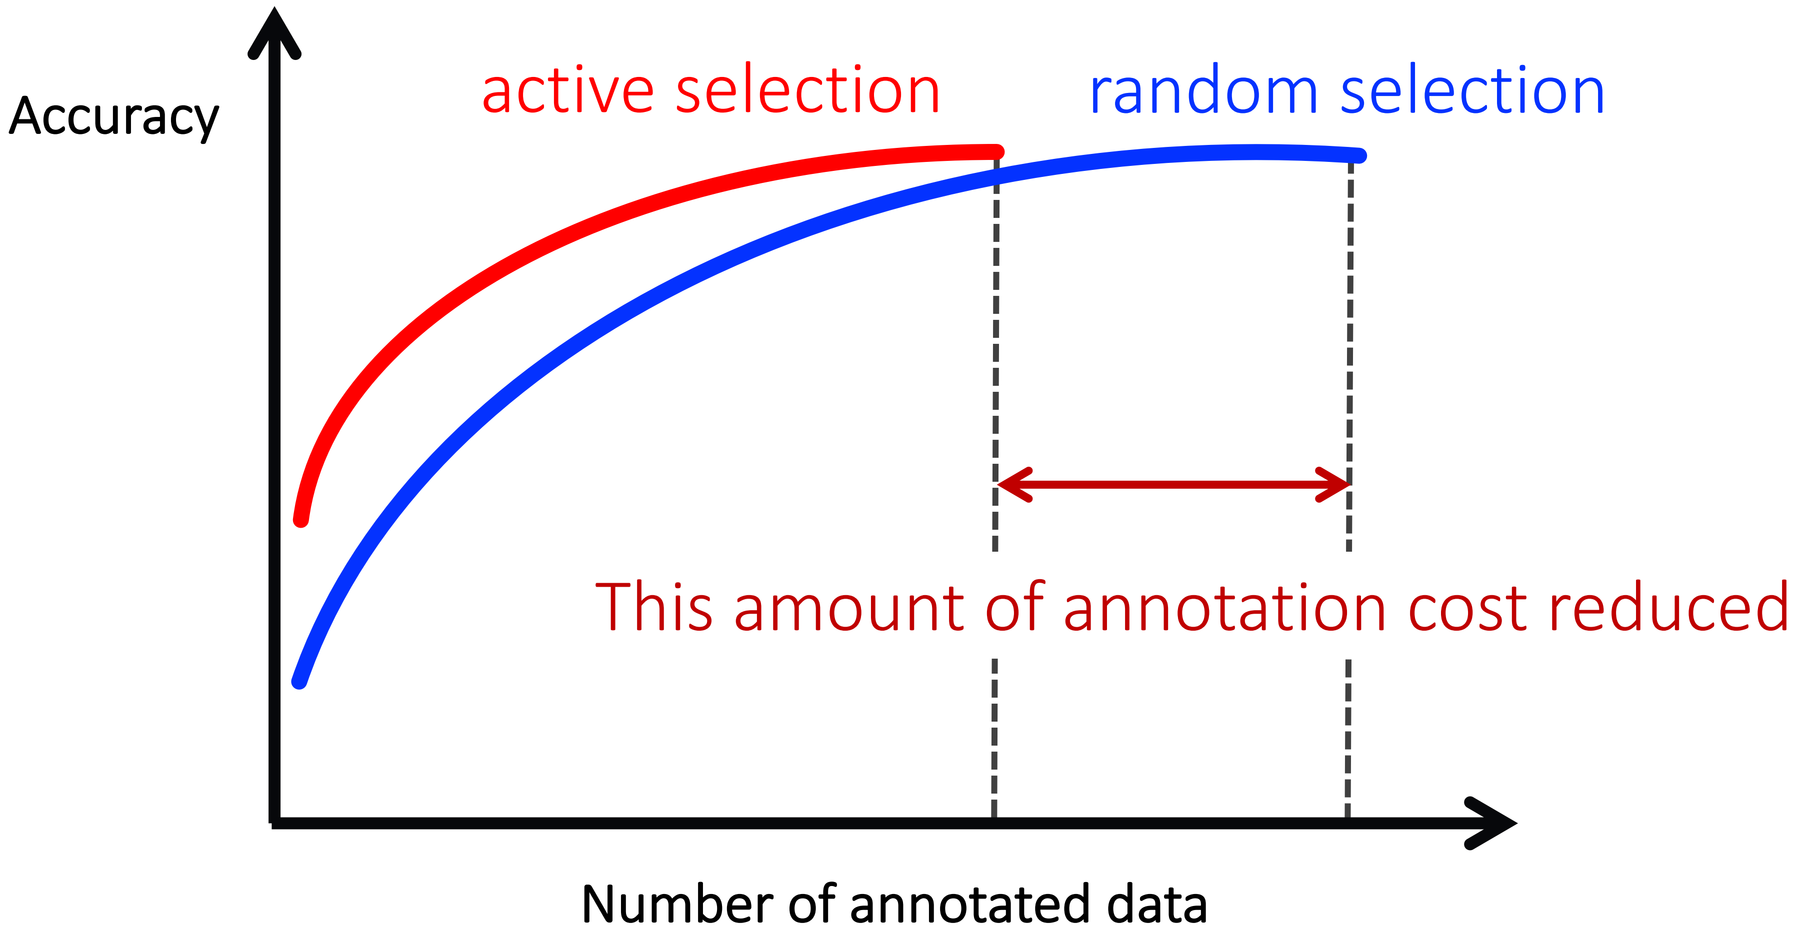

Aim #1: Acquiring necessary annotation efficiently from human experts

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

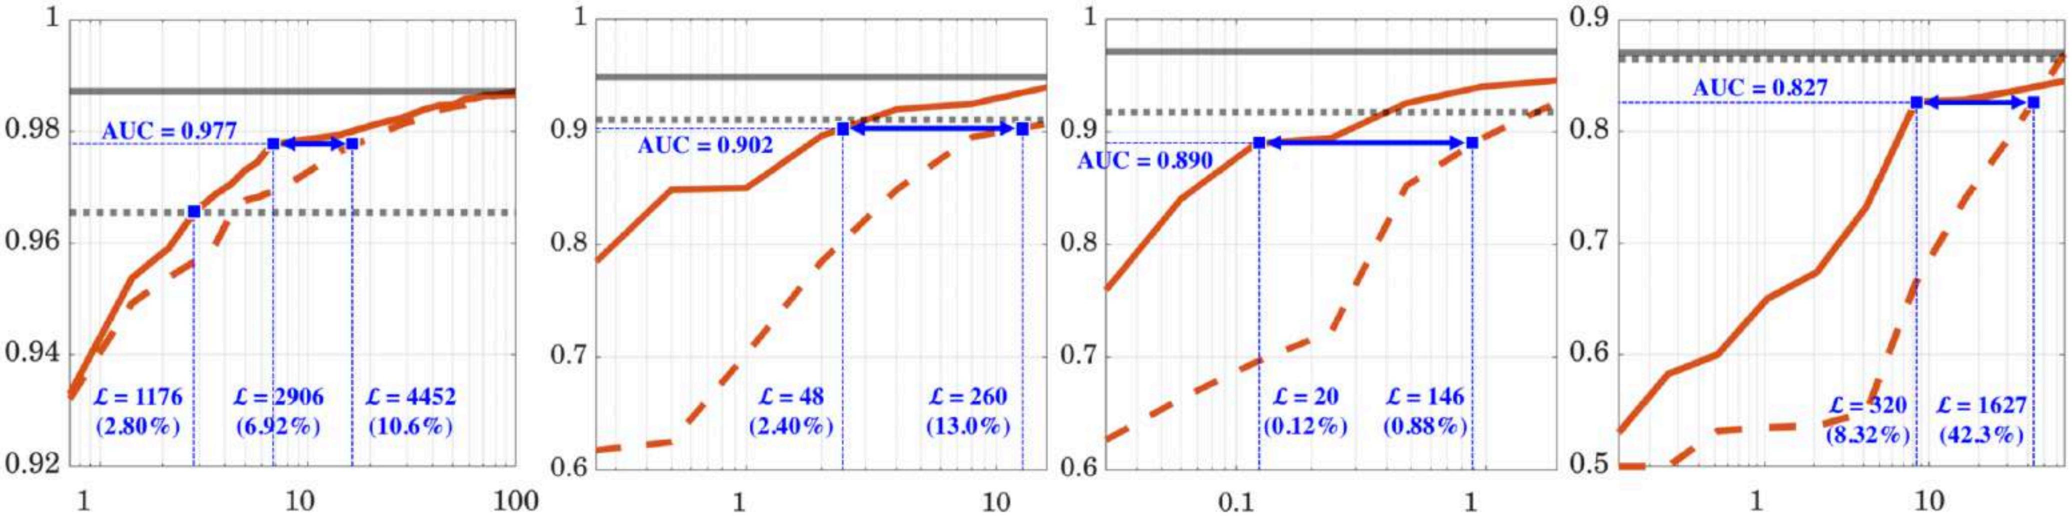

Places: Scene Classification

Colonoscopy Frame Classification

Polyp Detection

Pulmonary Embolism Detection

1. Zhou, Zongwei, et al. "Integrating active learning and transfer learning for carotid intima-media thickness video interpretation." Journal of digital imaging 32.2 (2019): 290-299.
2. Zhou, Zongwei, et al. "Active, Continual Fine Tuning of Convolutional Neural Networks for Reducing Annotation Efforts." arXiv preprint arXiv:1802.00912 (2018).
3. Zhou, Zongwei, et al. "Fine-tuning convolutional neural networks for biomedical image analysis: actively and incrementally." Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.

Aim #1: Acquiring necessary annotation efficiently from human experts

Proposal: Iteratively suggest important samples at the patient-level

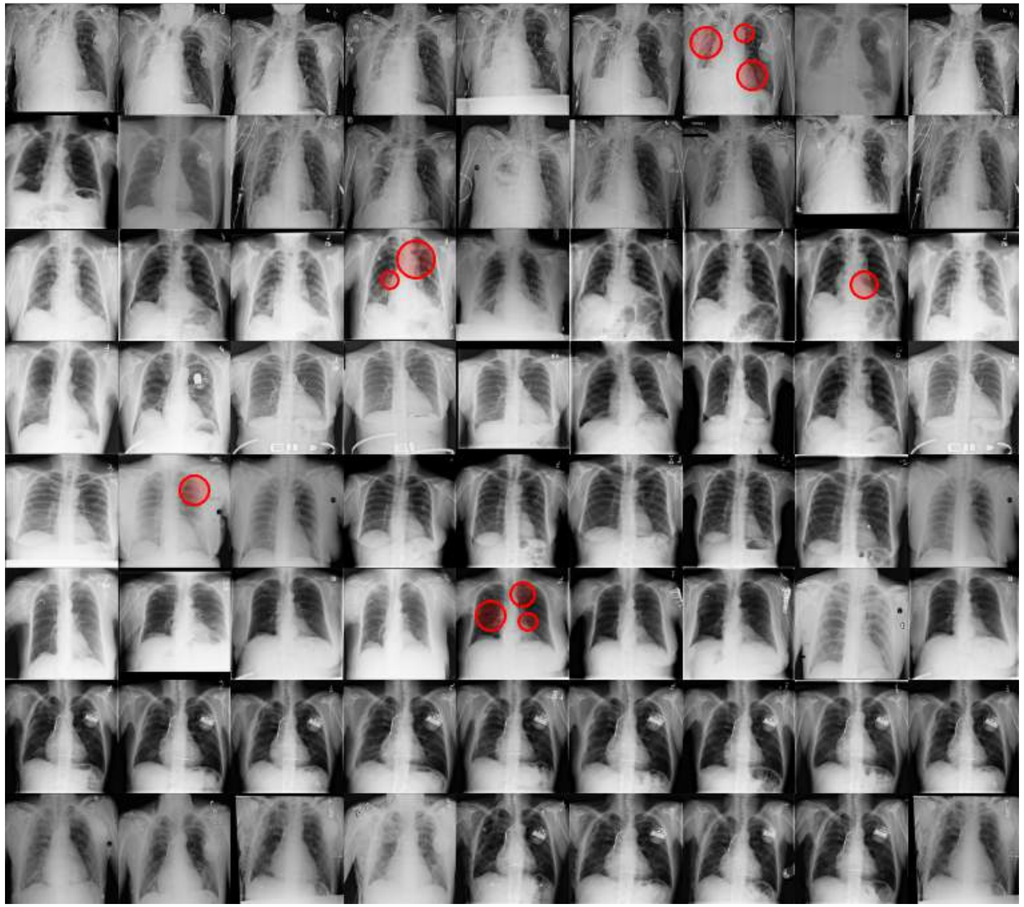

Lesion-level annotation

Drawbacks:

- Experts must annotate the same patient multiple times

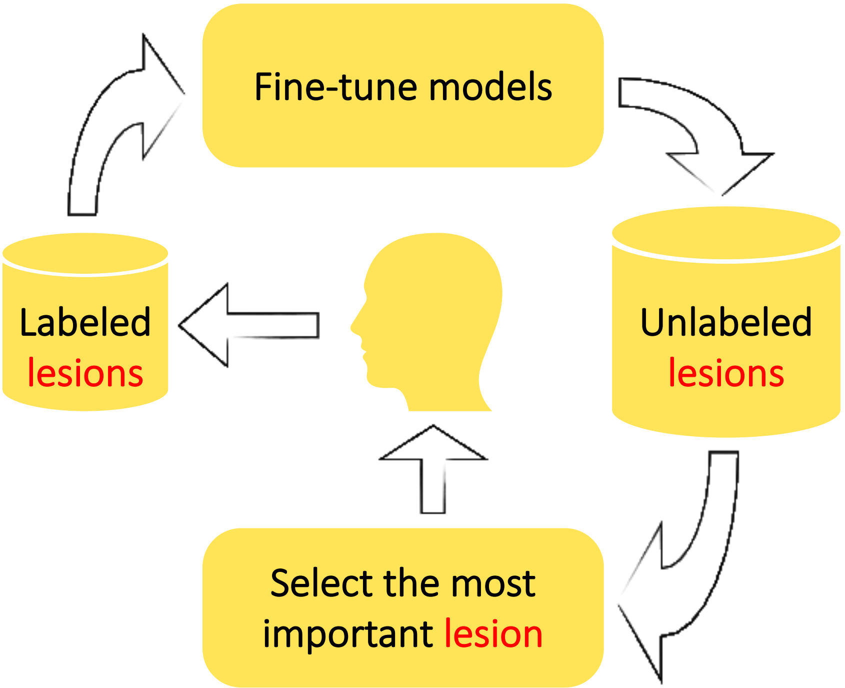

Aim #1: Acquiring necessary annotation efficiently from human experts

Proposal: Iteratively suggest important samples at the patient-level

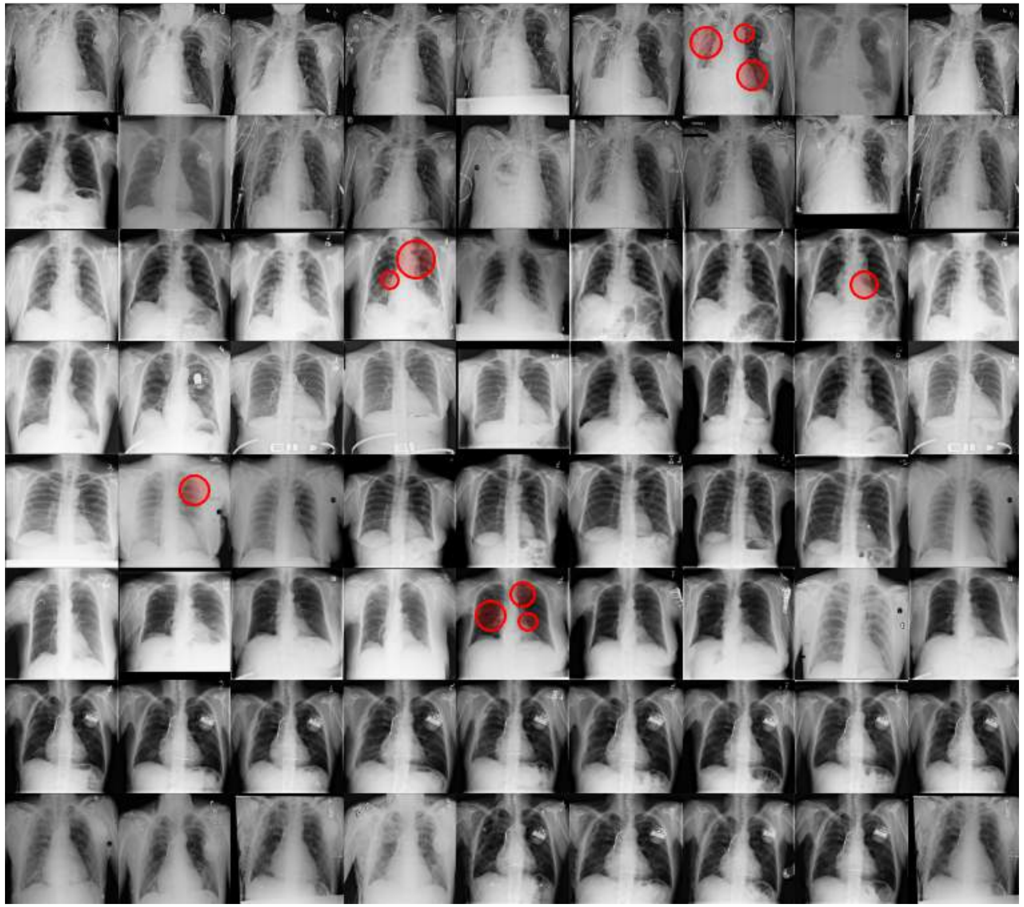

Lesion-level annotation

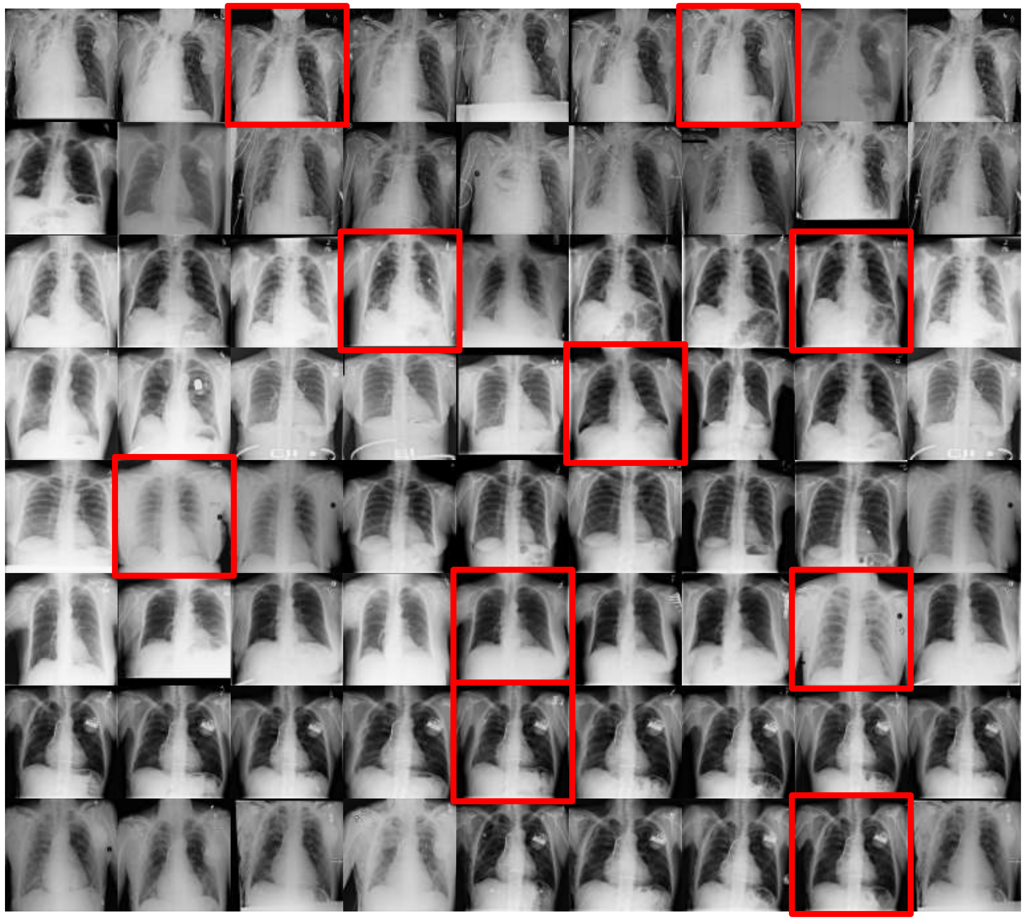

Patient-level annotation

Not All Data Is Created Equal

Featured Publications for Aim #1:

1. Z. Zhou, J. Shin, L. Zhang, S. Gurudu, M. Gotway, J. Liang, 2017. Fine-tuning Convolutional Neural Networks for Biomedical Image Analysis: Actively and Incrementally. CVPR'17, one of only five papers in biomedical imaging accepted by CVPR’17.
2. Z. Zhou, J. Shin, R. Feng, R. Hurst, C. Kendall, J. Liang, 2019. Integrating Active Learning and Transfer Learning for Carotid Intima-Media Thickness Video Interpretation. Journal of Digital Imaging.
3. Z. Zhou, J. Shin, S. Gurudu, M. Gotway, J. Liang, 2020. Active, Continual Fine Tuning of Convolutional Neural Networks for Reducing Annotation Efforts. Submitted to Medical Image Analysis.

Not All Data Is Created Equal

Clinical Impacts of Aim #1:

1. The continual learning capability of of deep models encourages data, label, and model reuse.
2. An efficient “human-in-the-loop” procedure assists radiologists in quickly dismissing patients with negative results, therefore dramatically reducing the burden of annotation.
3. An instant on-line feedback process makes it possible for CAD systems to be self-learning and self-improving via continual fine-tuning.

Introduction

Objective

Aim #1

Aim #2

Aim #3

Summary

Aim #2: Utilizing existing annotation effectively from advanced architecture

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

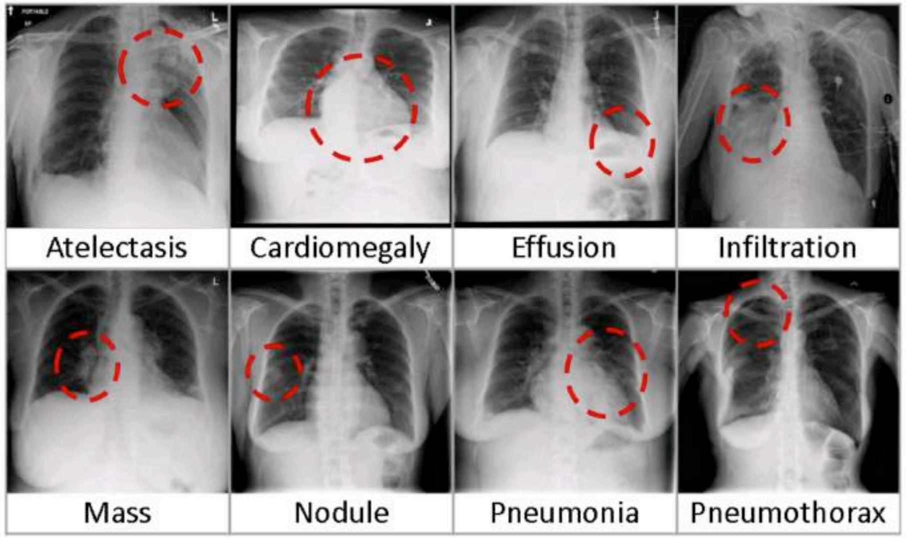

$ 1,000 annotation budget ☺

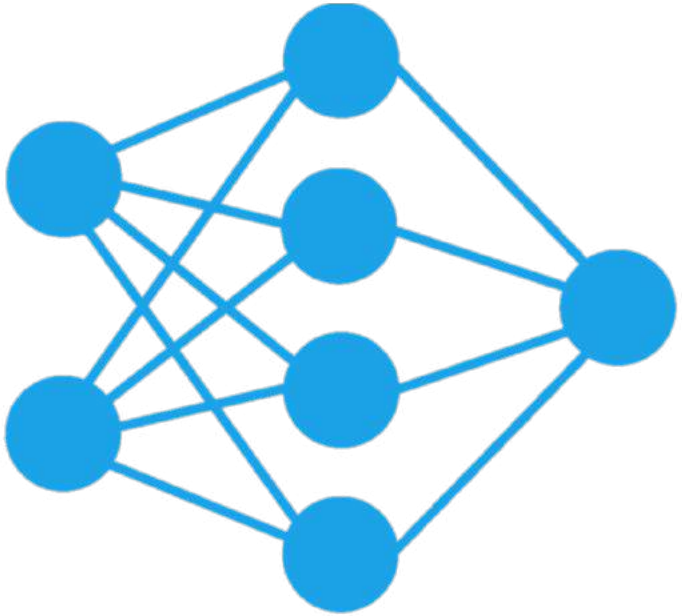

Aim #2: Utilizing existing annotation effectively from advanced architecture

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

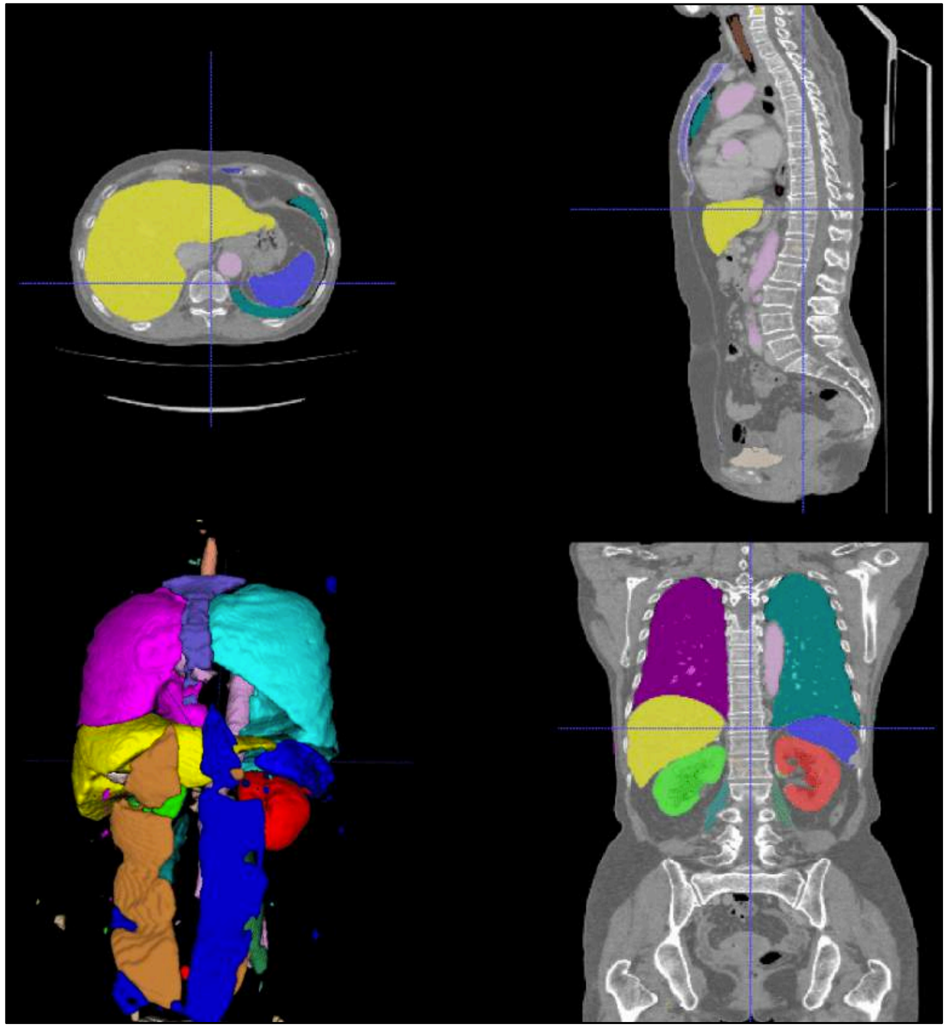

e.g., liver & lesion segmentation

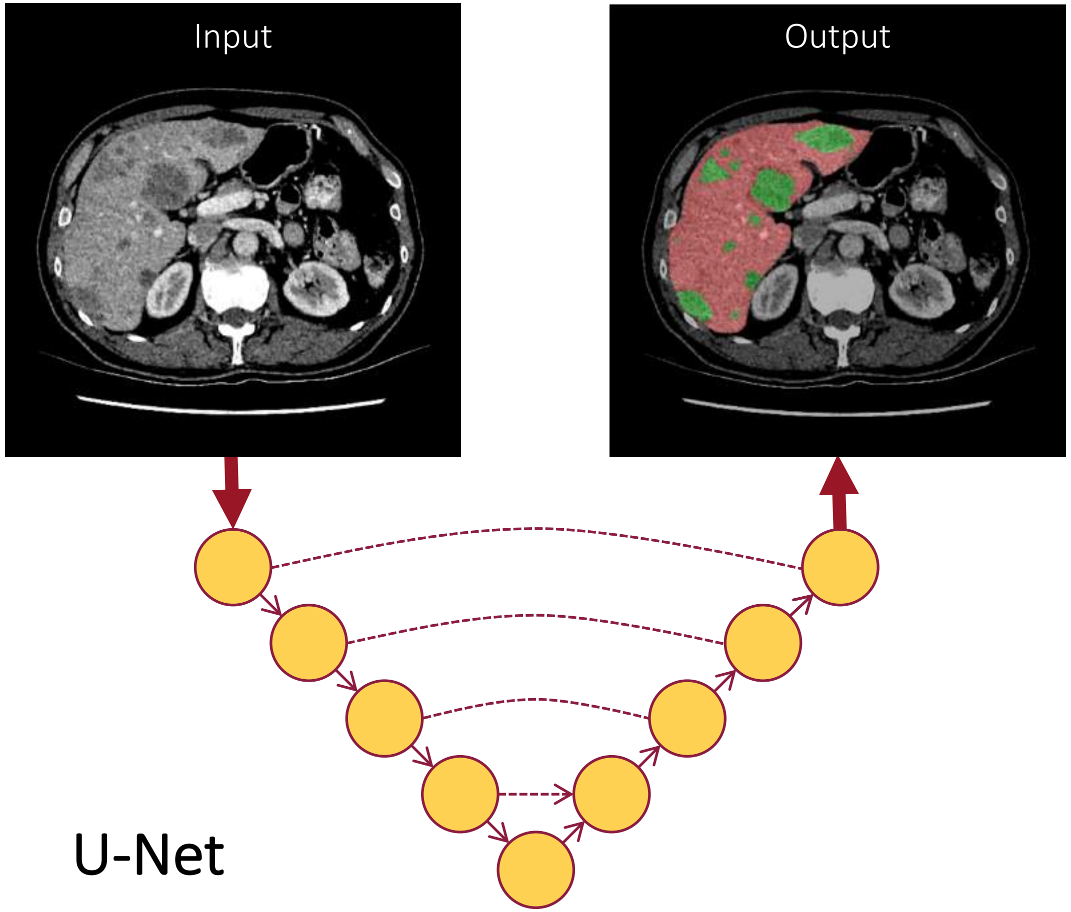

Aim #2: Utilizing existing annotation effectively from advanced architecture

Hypothesis: Multi-scale feature aggregation leads to powerful models

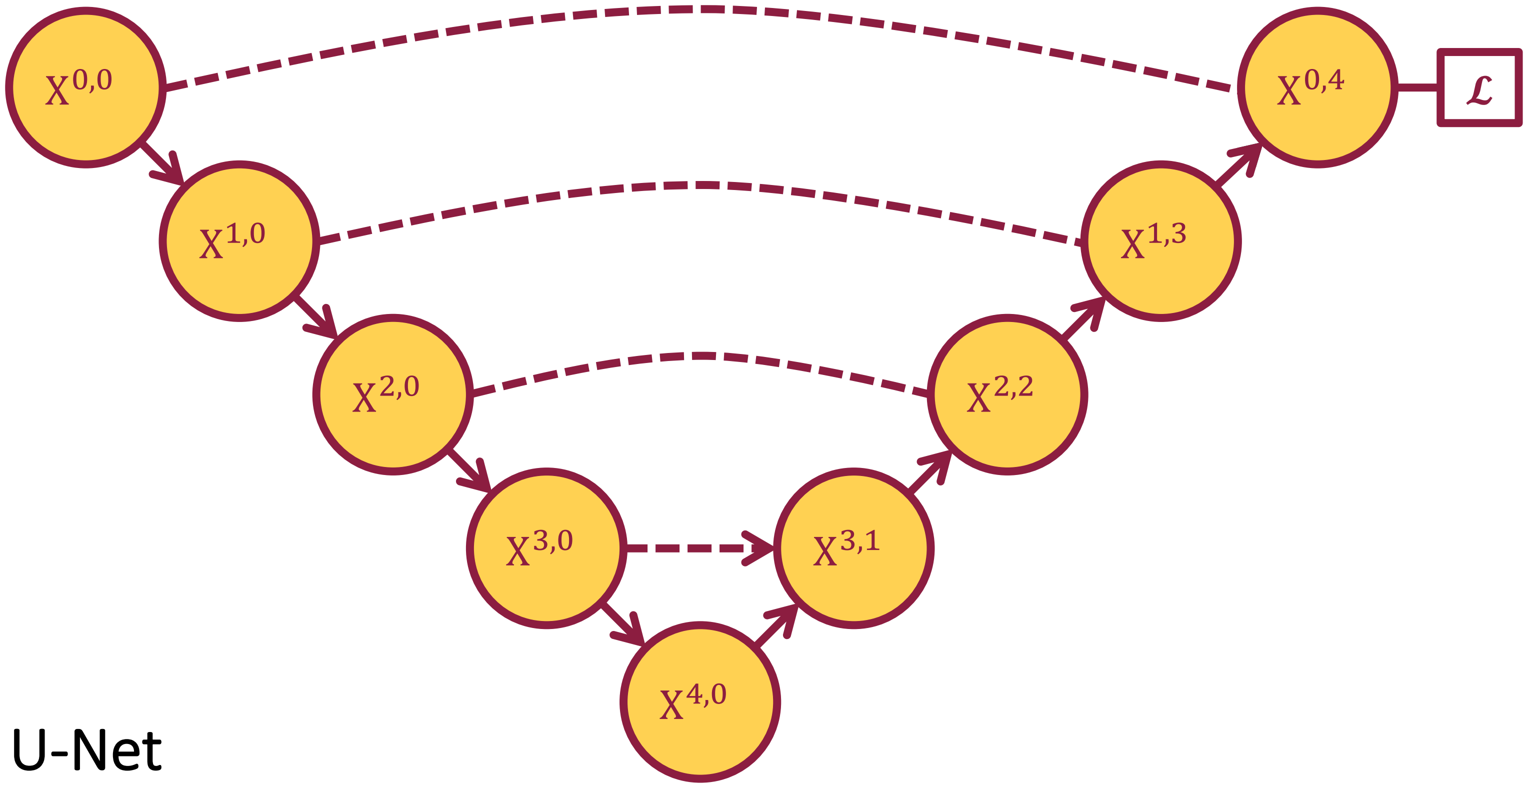

U-Net

1. Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. "U-net: Convolutional networks for biomedical image segmentation." International Conference on Medical image computing and computer-assisted intervention. Springer, Cham, 2015.

Aim #2: Utilizing existing annotation effectively from advanced architecture

Approach: Redesigned skip connections aggregate multi-scale features

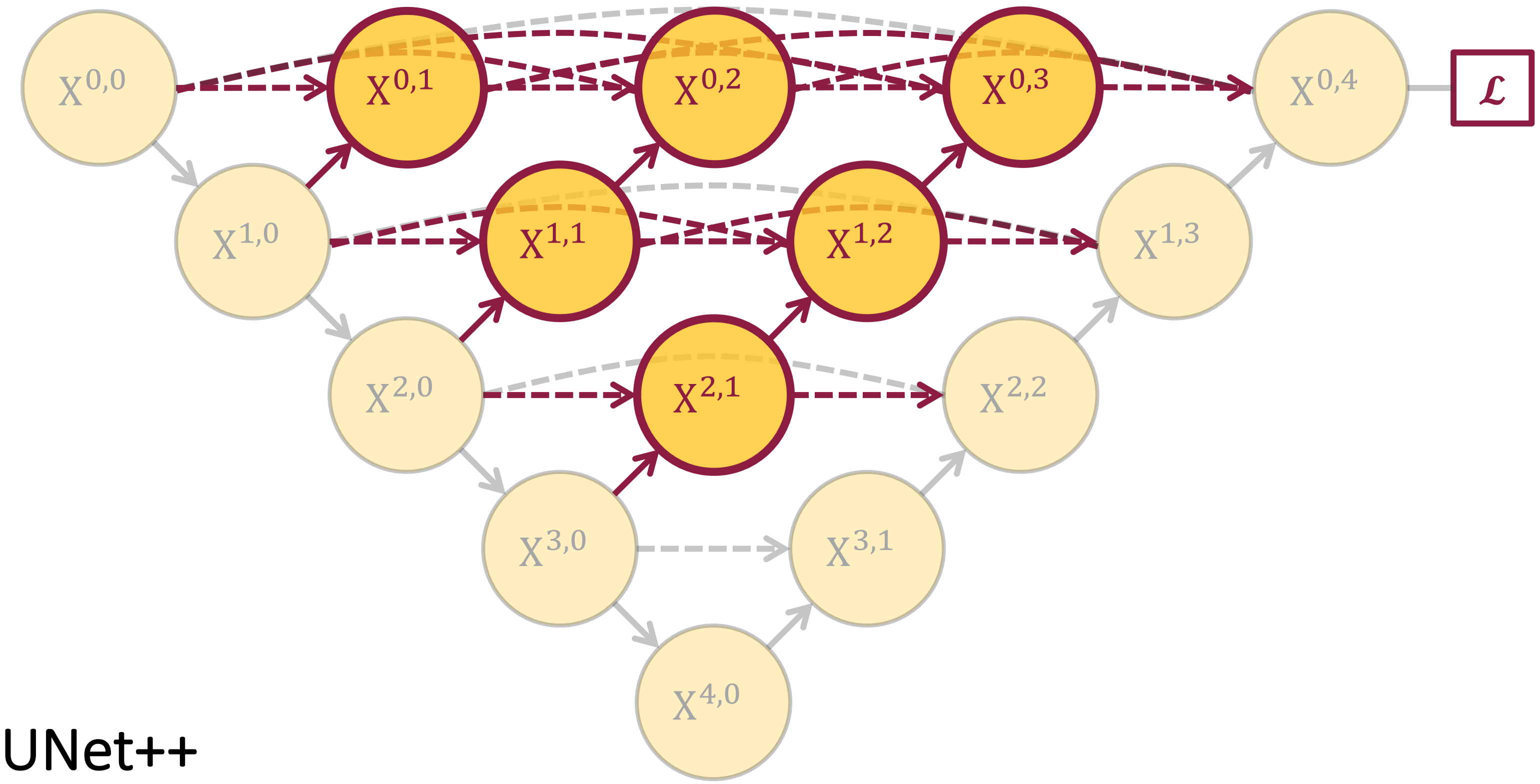

UNet++

1. Zhou, Zongwei, et al. "Unet++: A nested u-net architecture for medical image segmentation." Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. Springer, Cham, 2018. 3-11.

Aim #2: Utilizing existing annotation effectively from advanced architecture

Approach: Deep supervision enables a higher segmentation accuracy

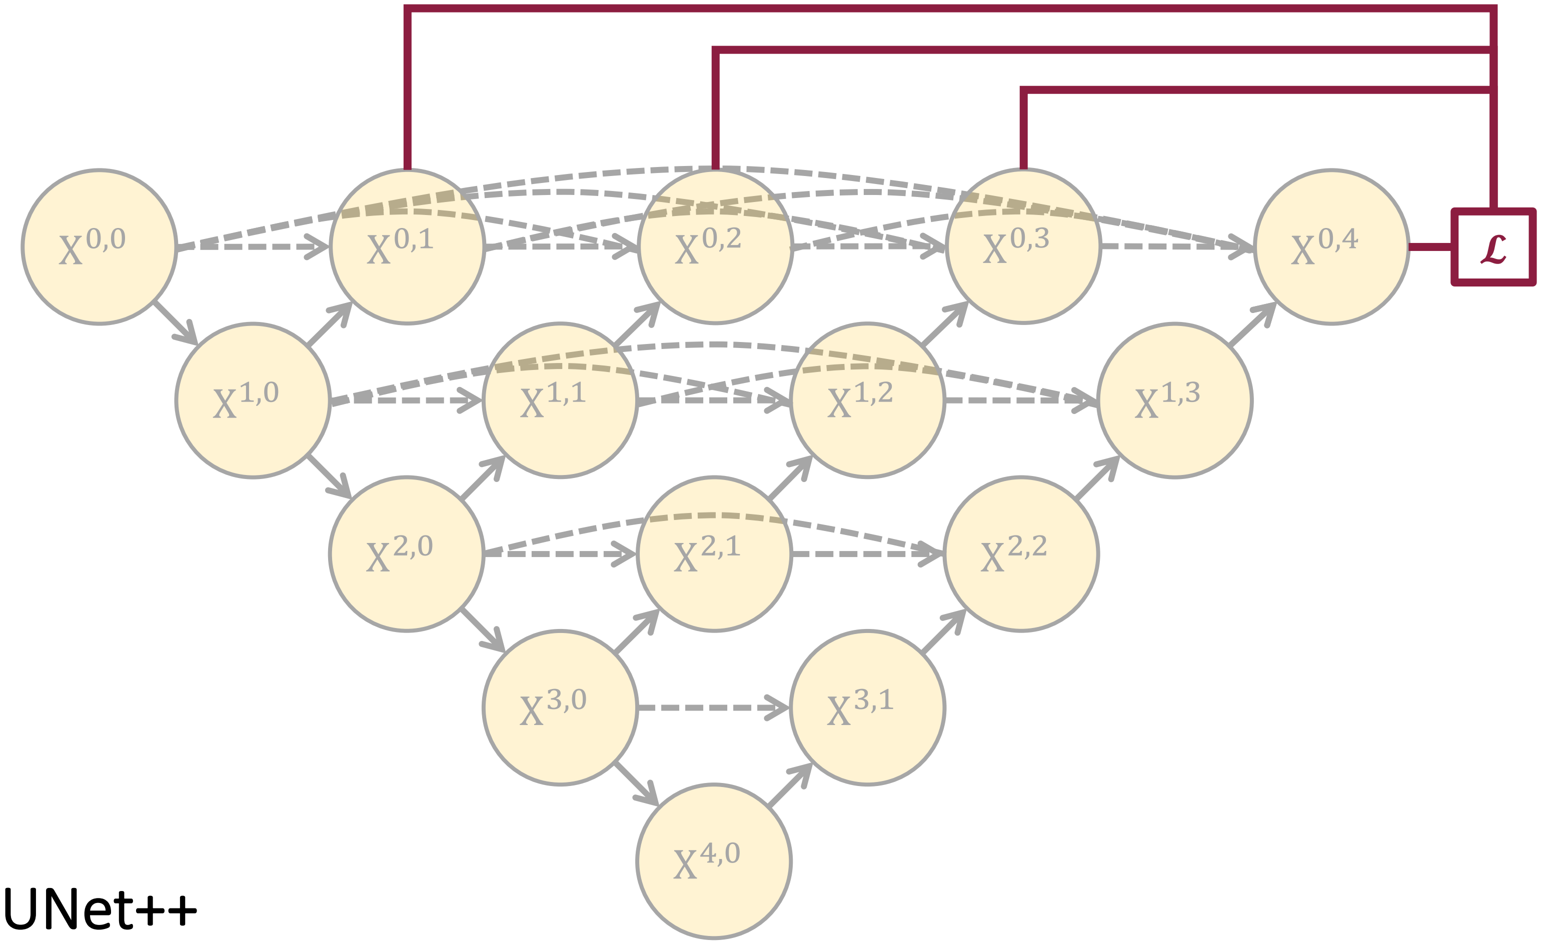

UNet++

1. Zhou, Zongwei, et al. "Unet++: A nested u-net architecture for medical image segmentation." Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. Springer, Cham, 2018. 3-11.

Aim #2: Utilizing existing annotation effectively from advanced architecture

Approach: Deep supervision enables a higher segmentation accuracy

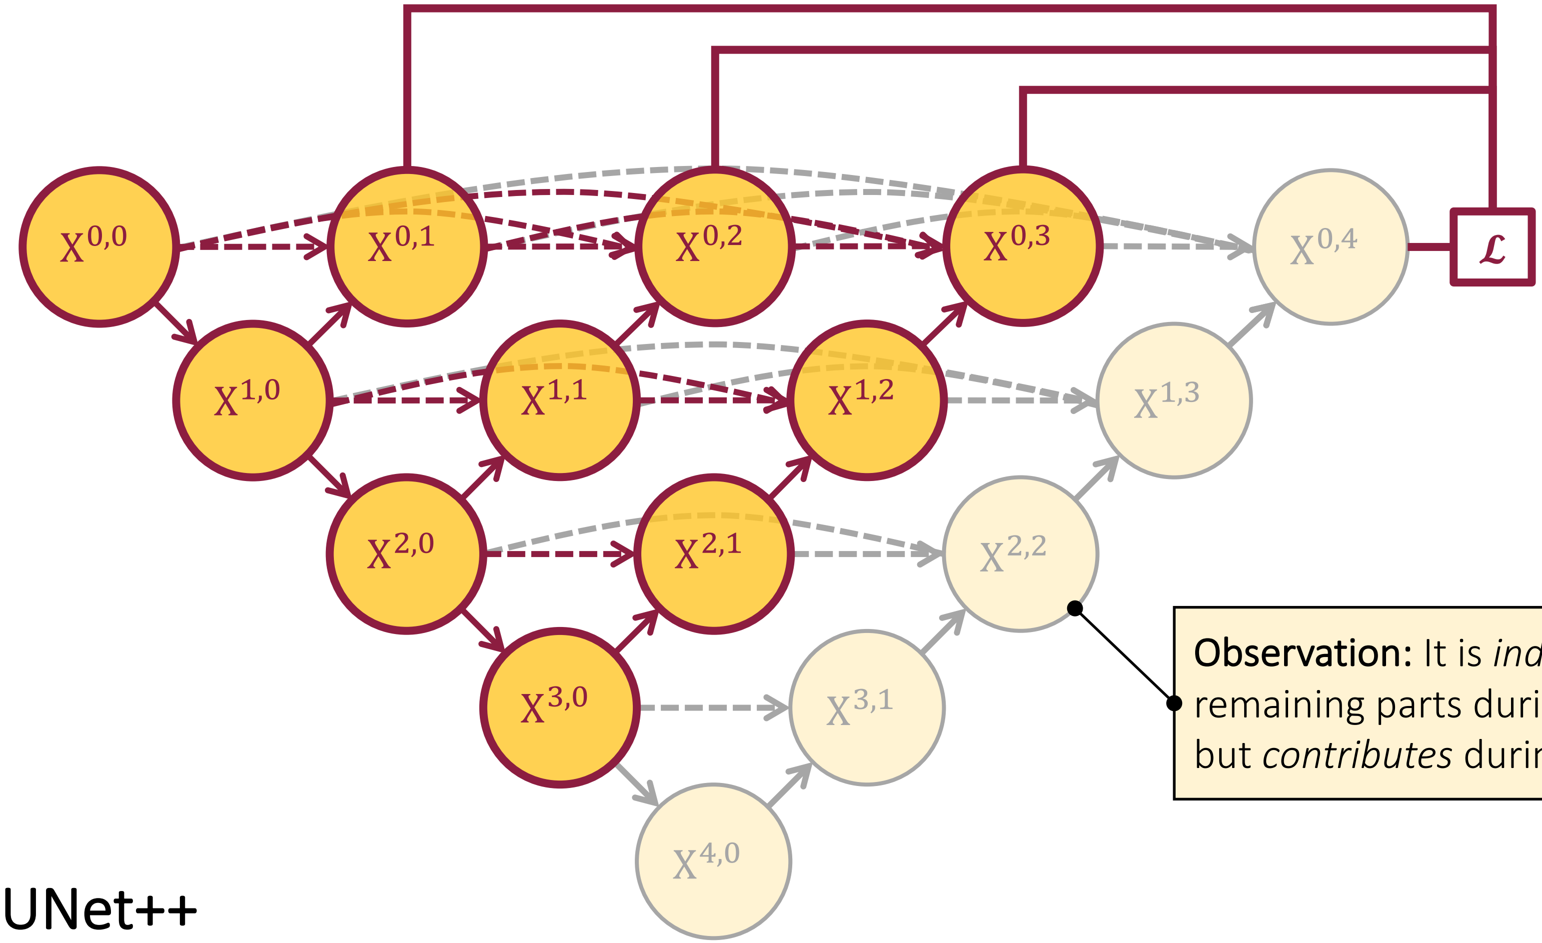

Observation: It is *independent* with the remaining parts during testing phase, but *contributes* during training phase.

UNet++

1. Zhou, Zongwei, et al. "Unet++: A nested u-net architecture for medical image segmentation." Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. Springer, Cham, 2018. 3-11.

Aim #2: Utilizing existing annotation effectively from advanced architecture

Contribution: UNet++ significantly improves disease/organ segmentation

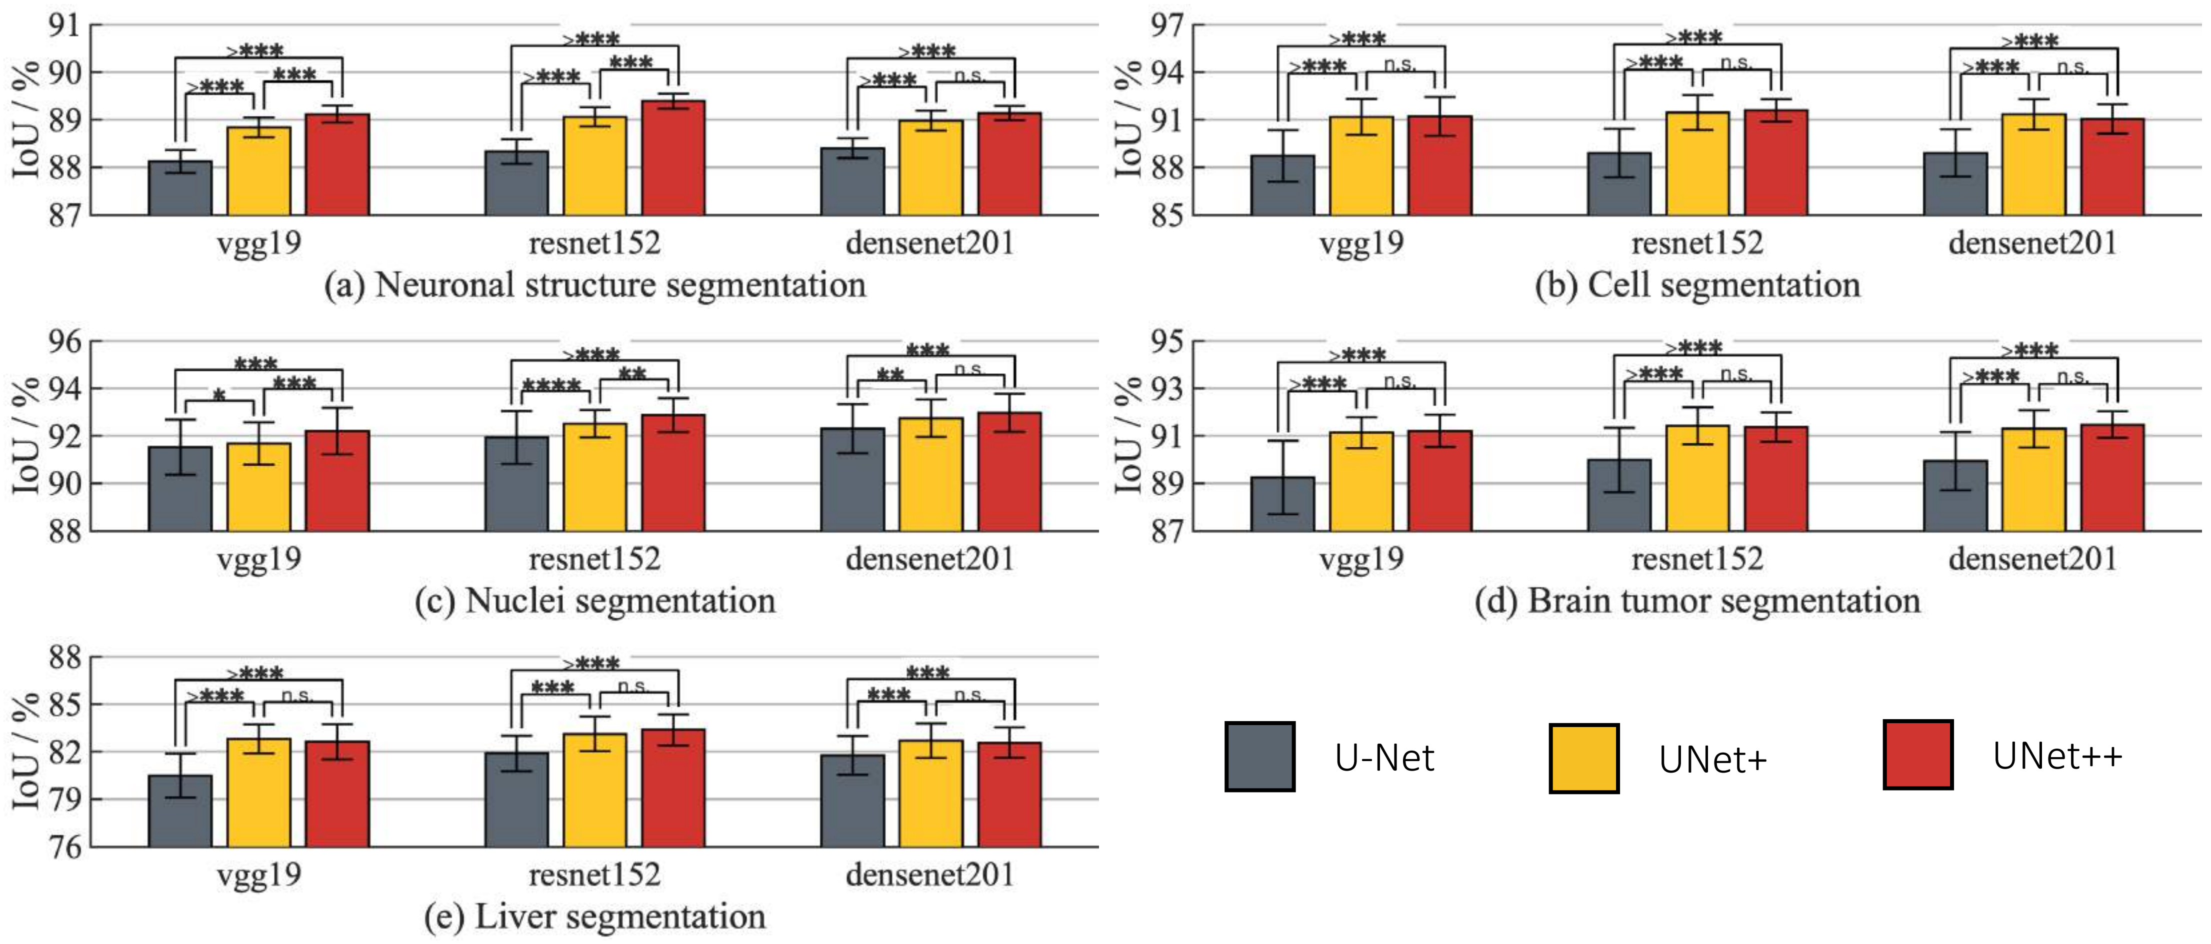

1. Zhou, Zongwei, et al. "Unet++: A nested u-net architecture for medical image segmentation." Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. Springer, Cham, 2018. 3-11.
2. Zhou, Zongwei, et al. "Unet++: Redesigning skip connections to exploit multiscale features in image segmentation." IEEE transactions on medical imaging 39.6 (2019): 1856-1867.

Introduction

Objective

Aim #1

Aim #2

Aim #3

Summary

Aim #2: Utilizing existing annotation effectively from advanced architecture

Proposal: Optimize active learning by leveraging unique architectural design

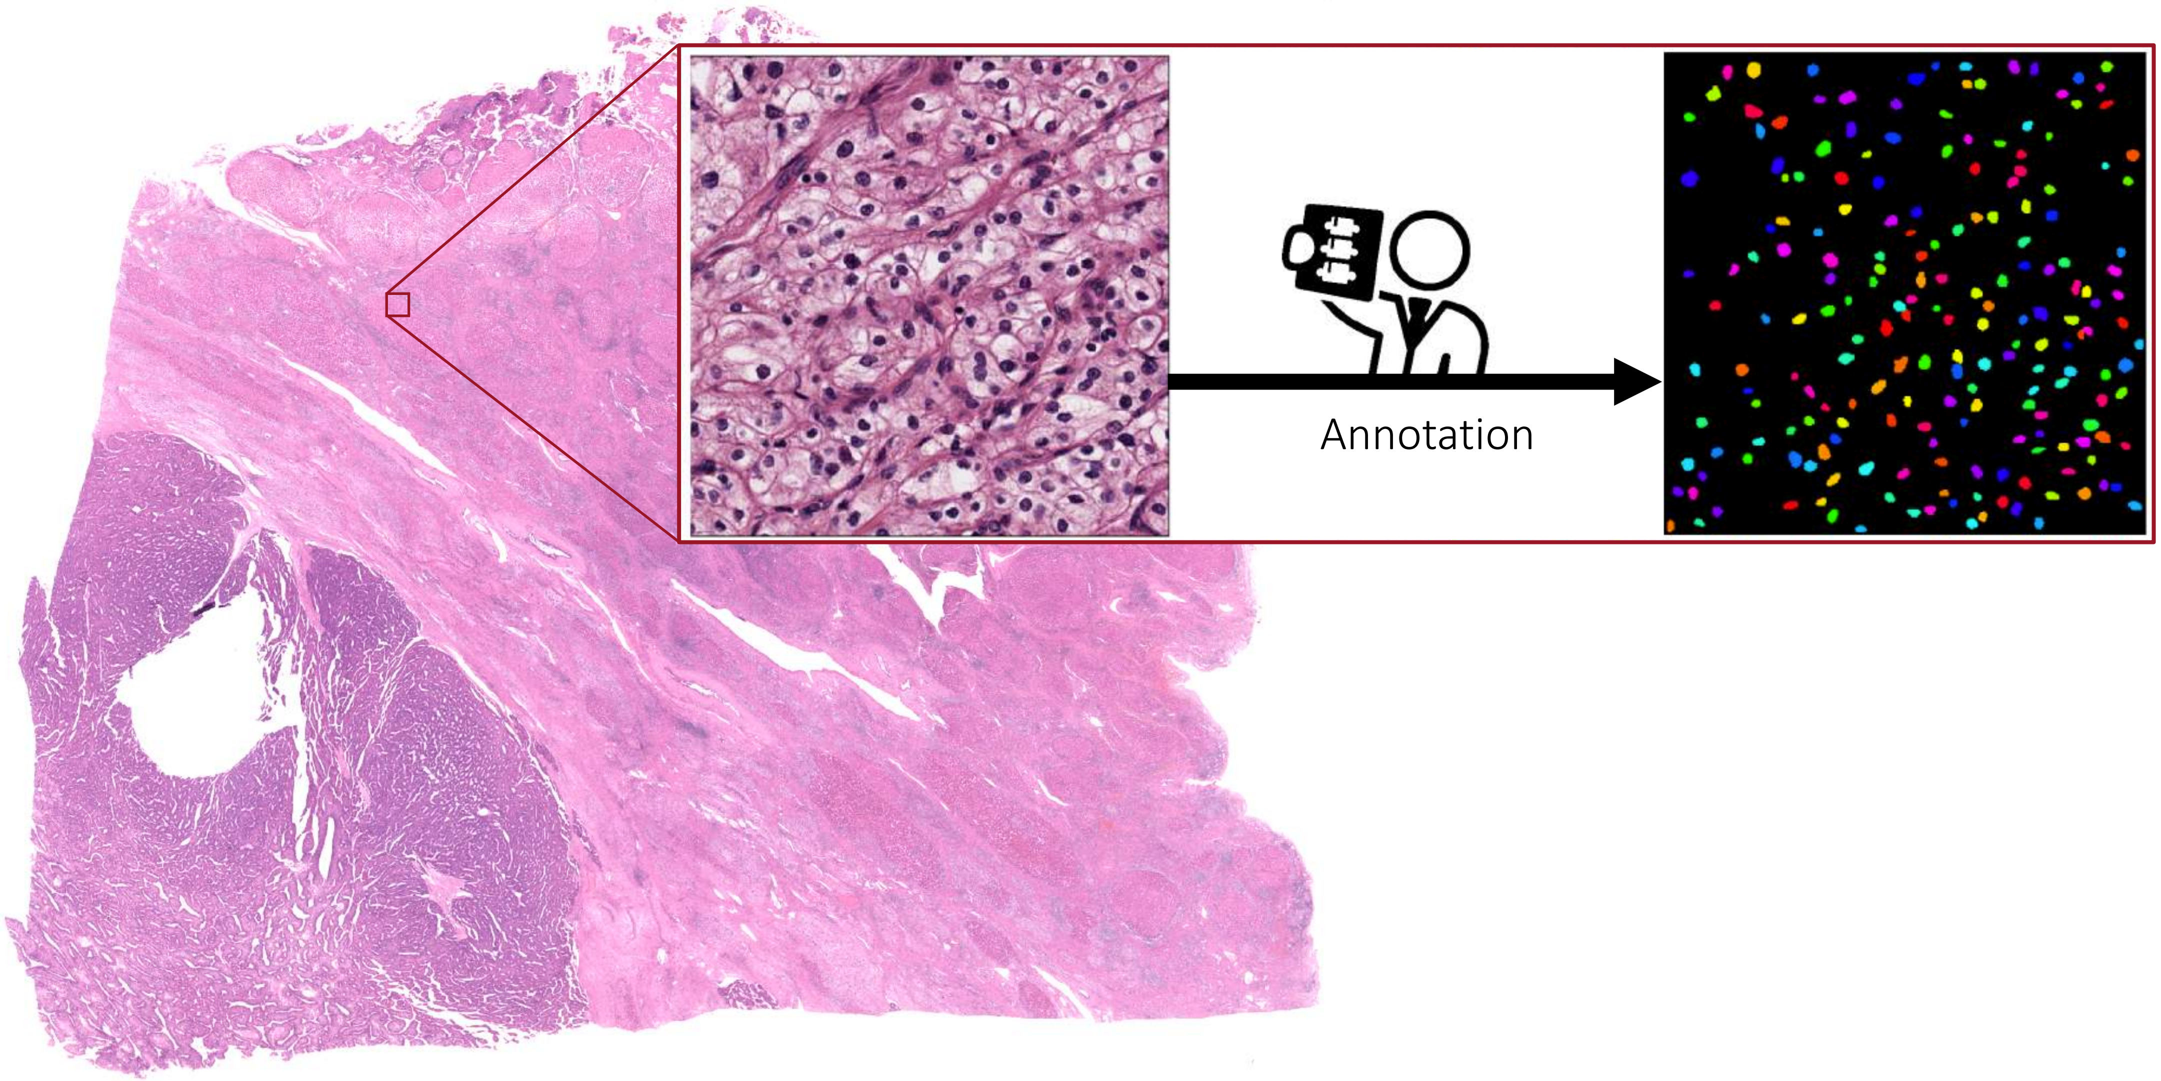

Aim #2: Utilizing existing annotation effectively from advanced architecture

Proposal: Optimize active learning by leveraging unique architectural design

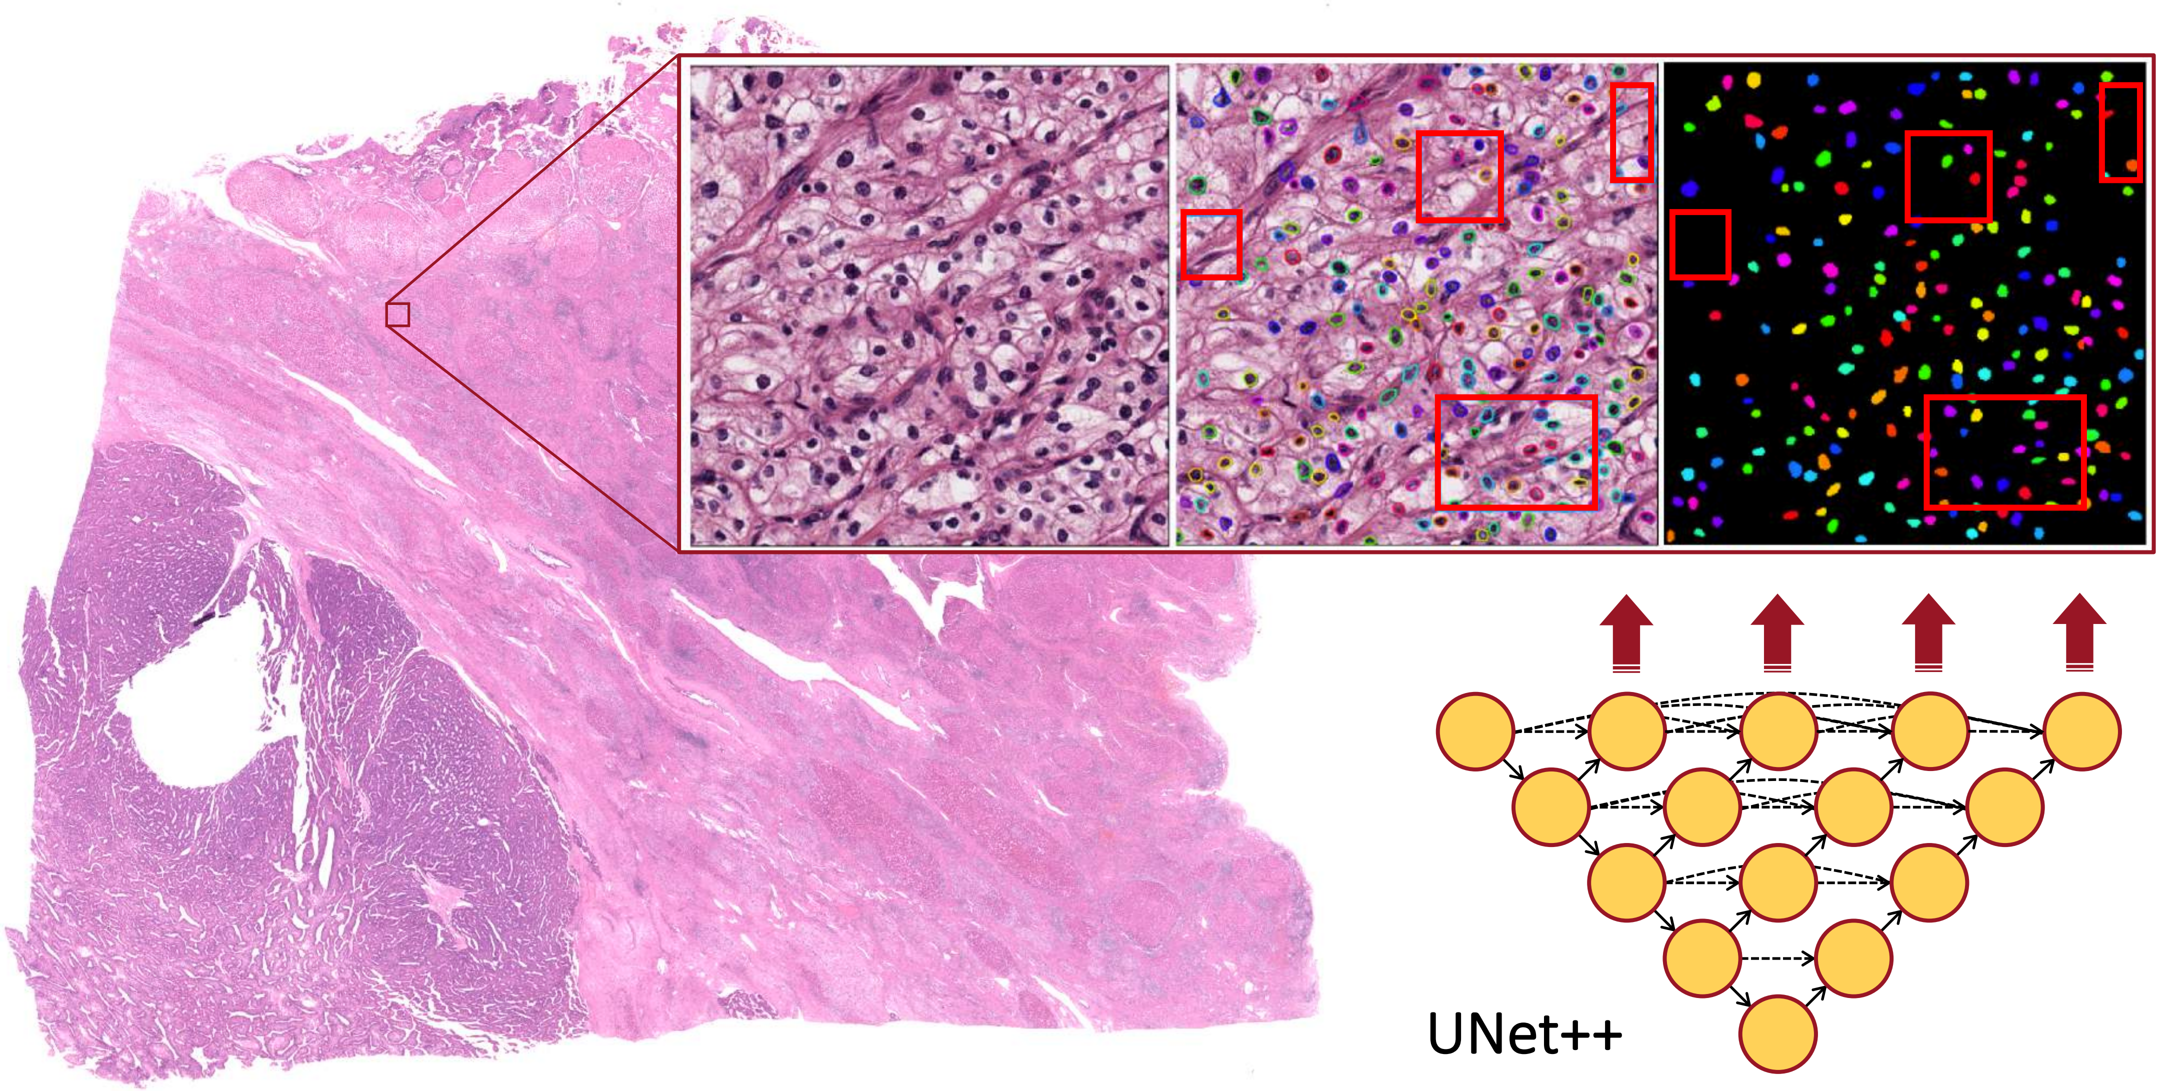

Intertwine the visual representation

Featured Publications for Aim #2:

1. Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, J. Liang, 2019. UNet++: Redesigning Skip Connections to Exploit Multi-Resolution Features in Image Segmentation. IEEE Transactions on Medical Imaging, IEEE TMI most popular articles.
2. Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, J. Liang, 2018. UNet++: A Nested U-Net Architecture for Medical Image Segmentation. DLMIA'18.

Intertwine the visual representation

Clinical Impacts of Aim #2:

1. Image segmentation can help compute clinically more accurate and desirable imaging bio-markers or precision measurement.
2. Model pruning has the potential to exert important impact on deploying computer-aided diagnosis (CAD) to mobile devices and ordinary desktop/laptop PCs in clinical practice.

Aim #3: Extracting generic knowledge directly from unannotated images

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

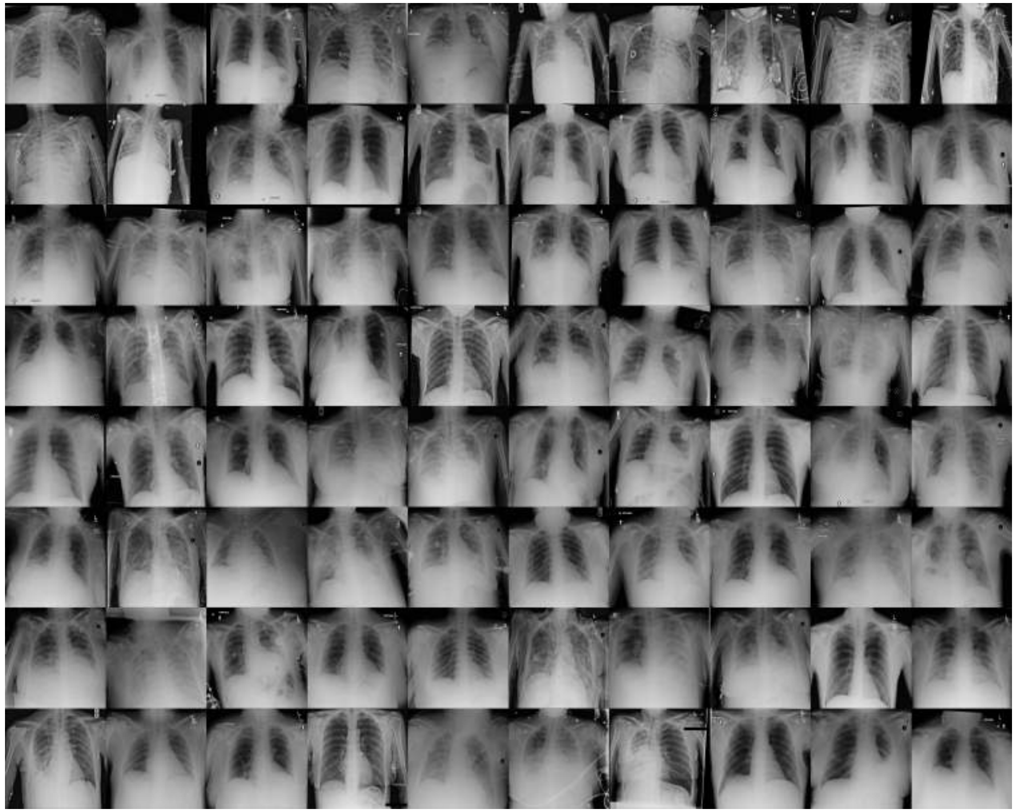

$ 1,000,000 annotation cost ☹

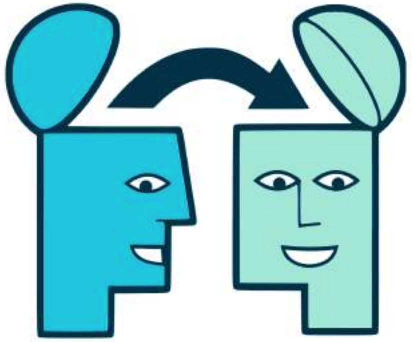

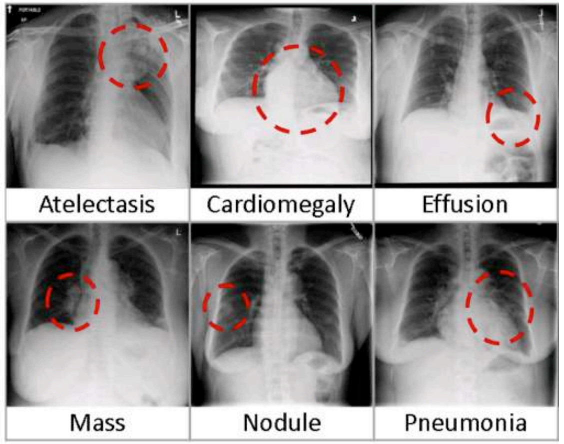

$ 1,000 annotation budget ☺

Aim #3: Extracting generic knowledge directly from unannotated images

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

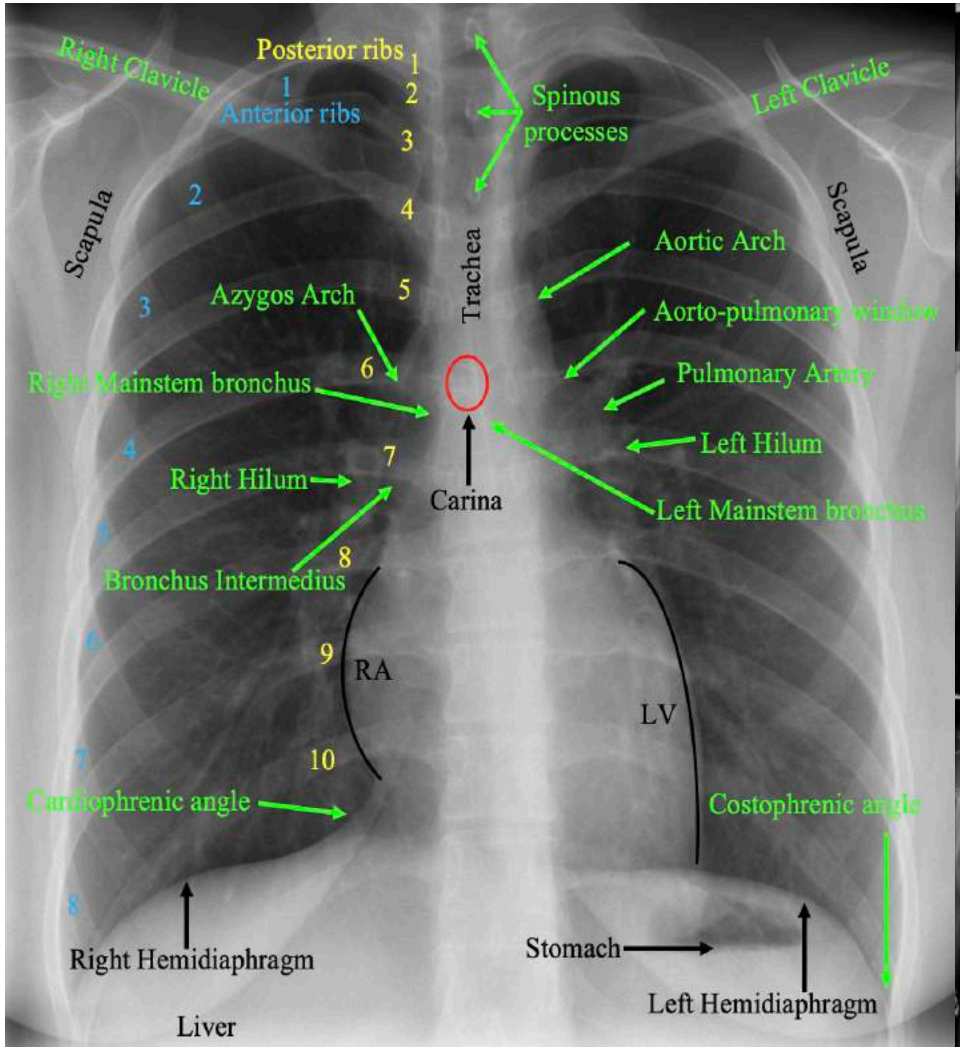

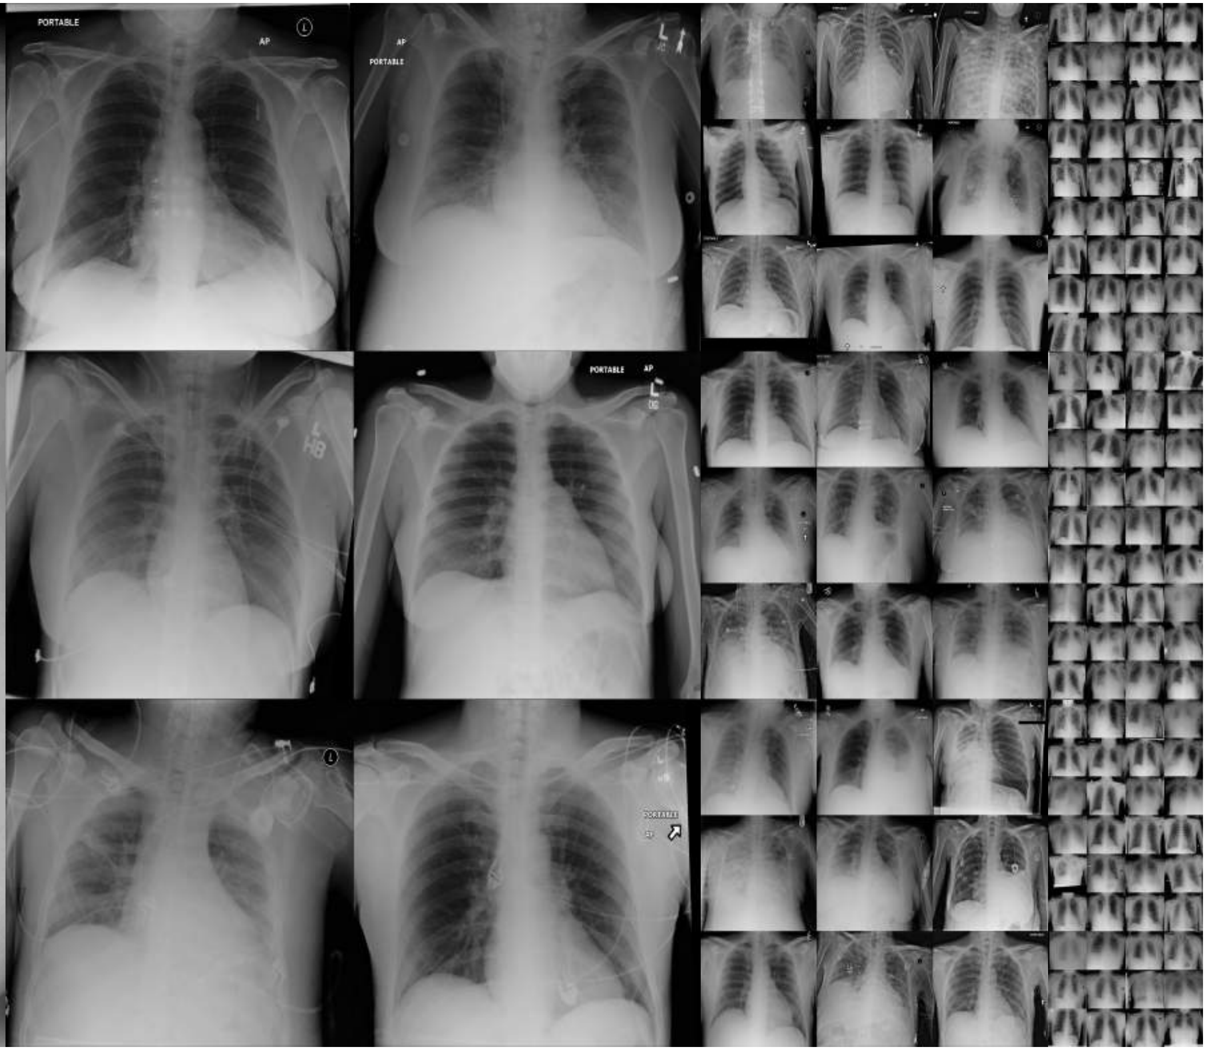

Aim #3: Extracting generic knowledge directly from unannotated images

Approach: Image restoration task helps model learn image representation

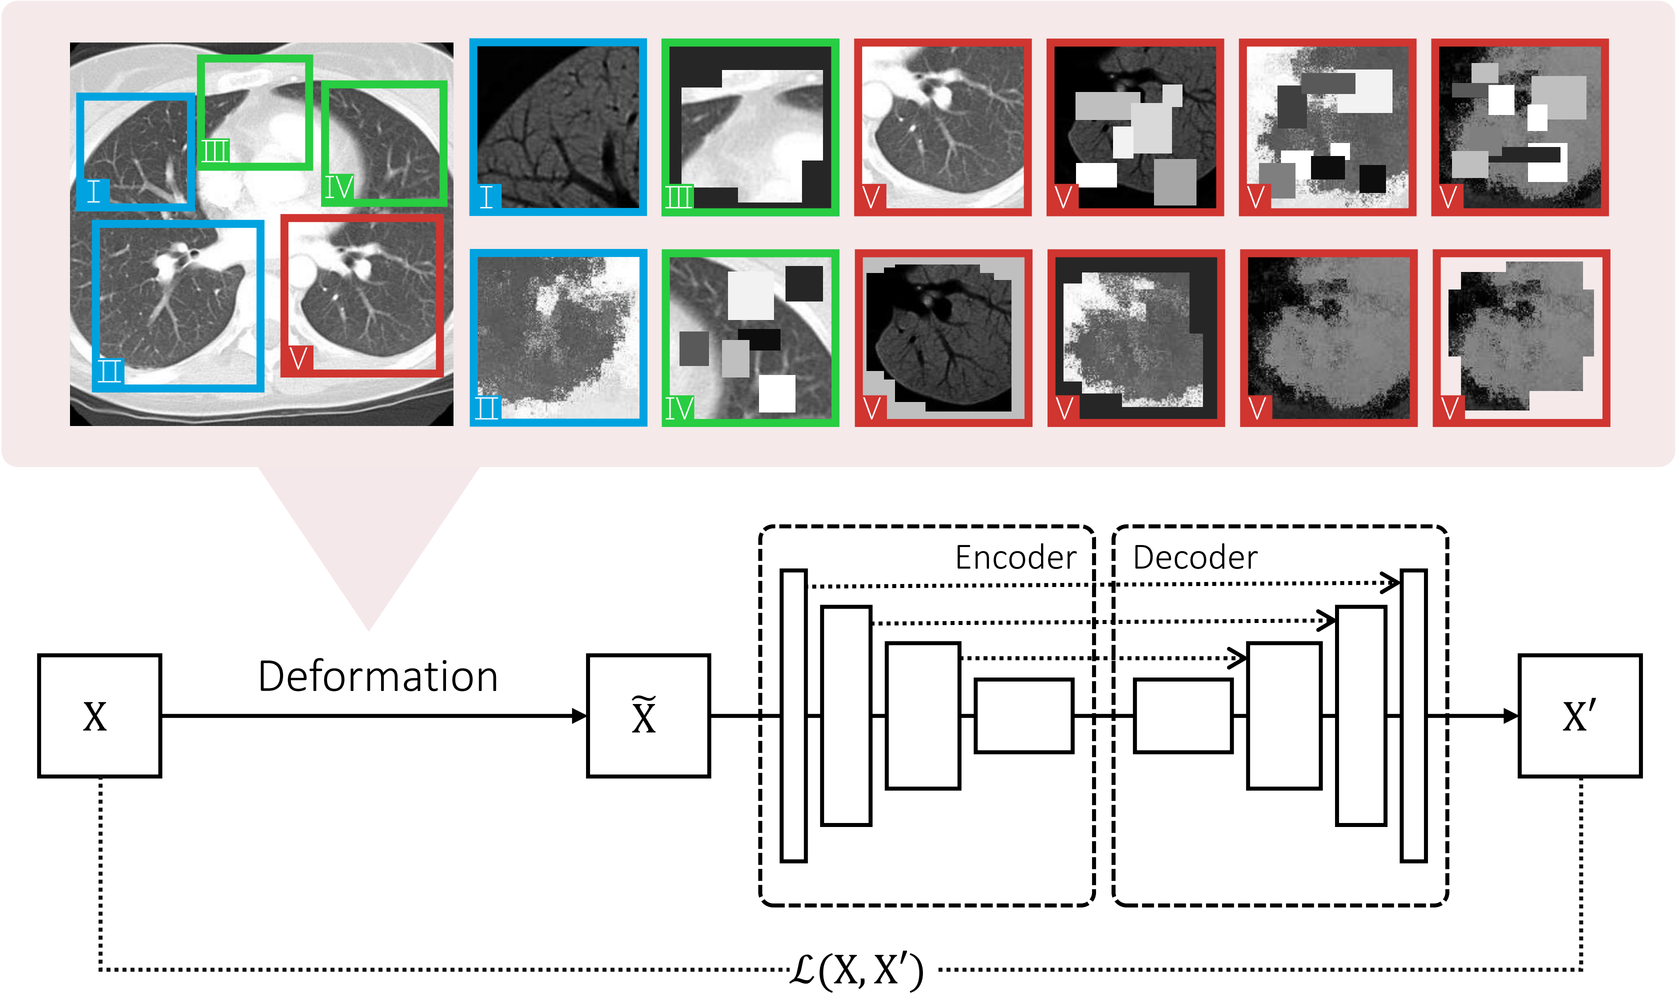

Aim #3: Extracting generic knowledge directly from unannotated images

Approach: Learning from multiple perspectives leads to robust models

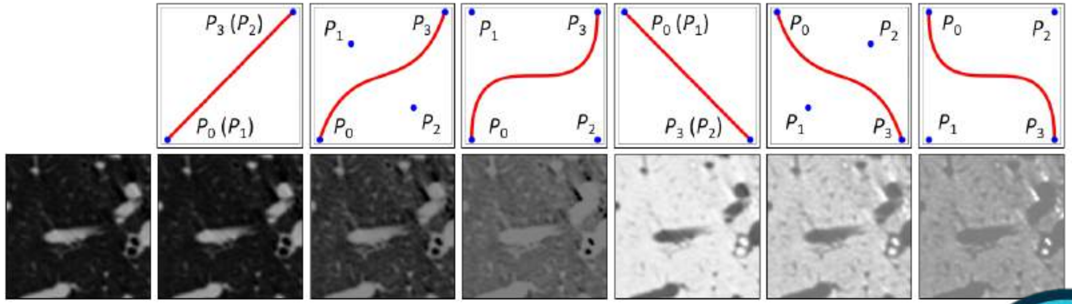

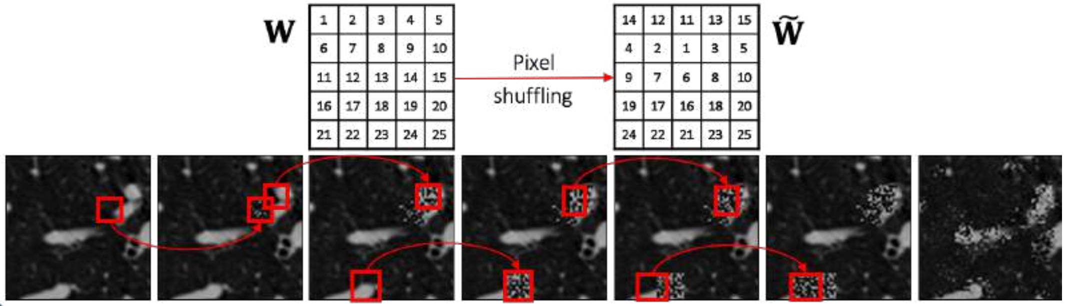

1. Learning organ appearance via **non-linear transformation**

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

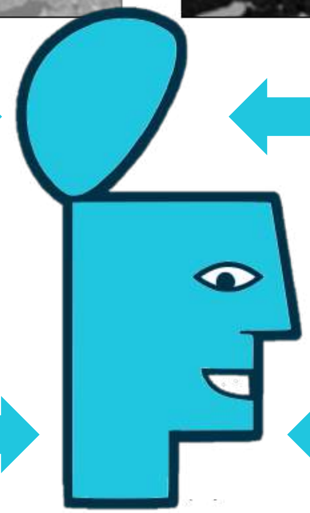

3. Learning organ spatial layout and global geometry via **outer-cutout**

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

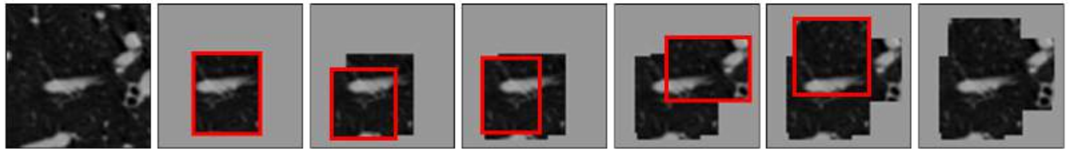

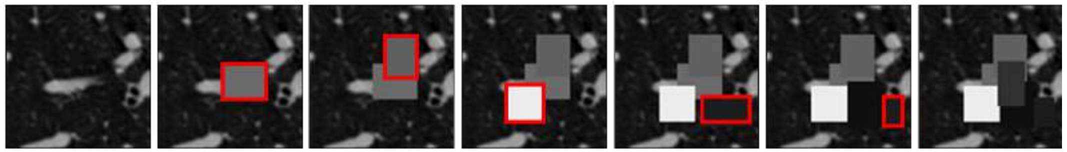

Introduction

Objective

Aim #1

Aim #2

Aim #3

Summary

Aim #3: Extracting generic knowledge directly from unannotated images

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

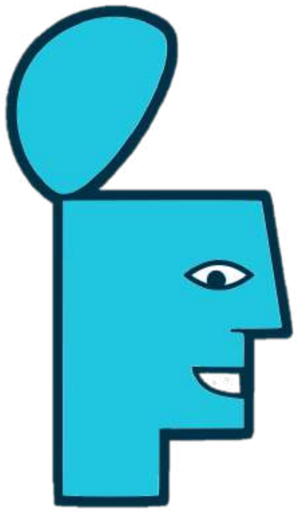

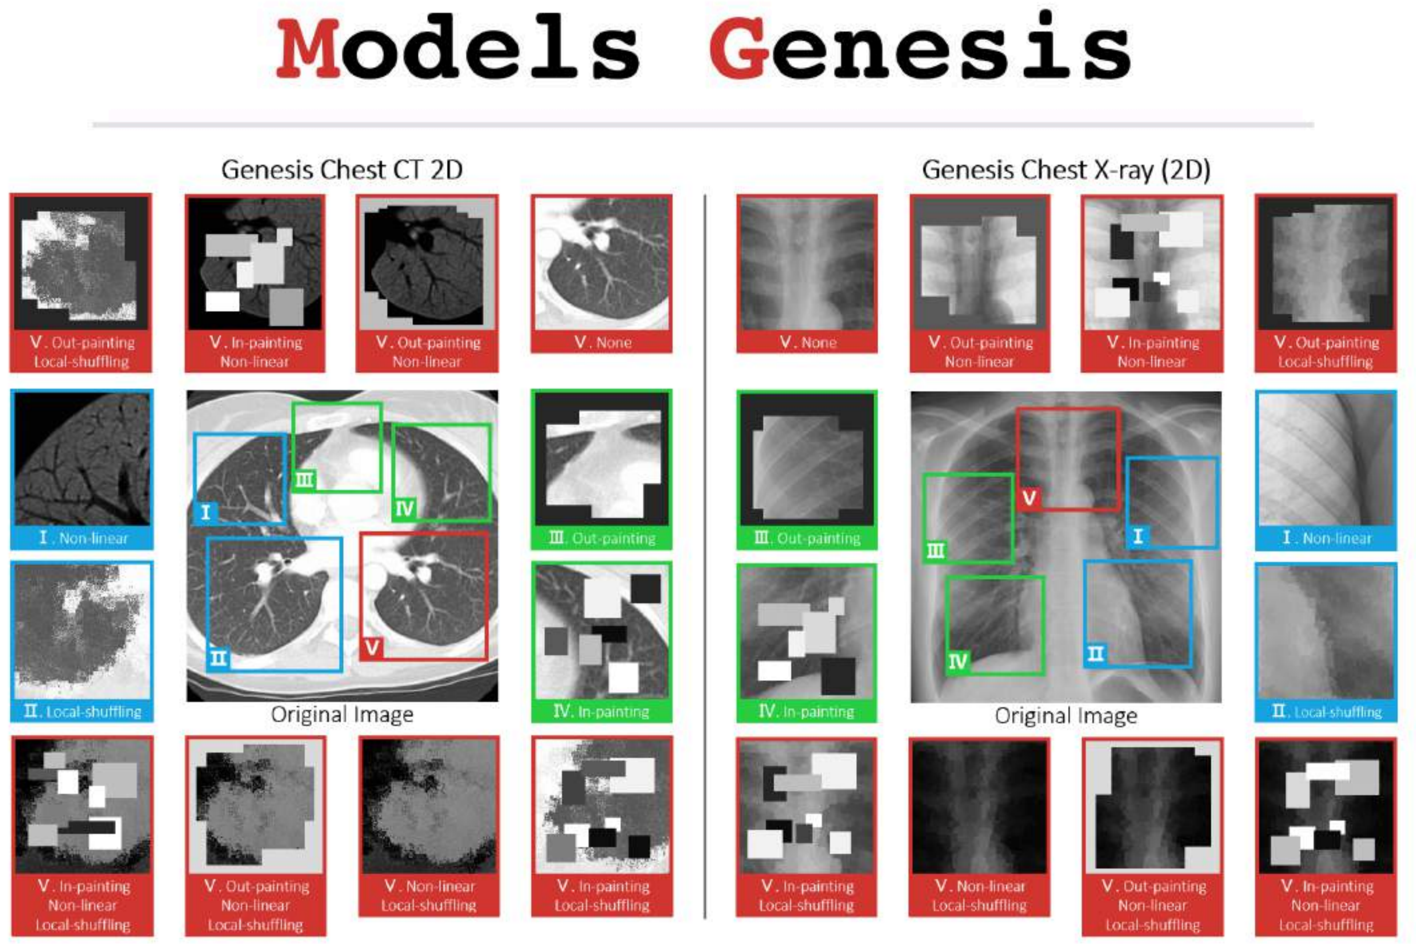

1. Zhou, Zongwei, et al. "Models genesis: Generic autodidactic models for 3d medical image analysis." International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, Cham, 2019.
2. Zhou, Zongwei, et al. "Models Genesis." arXiv preprint arXiv:2004.07882 (2020).

Aim #3: Extracting generic knowledge directly from unannotated images

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

Pre-training	Approach	Target tasks				
		NCC[1] (%)	NCS[2] (%)	ECC[3] (%)	LCS[4] (%)	BMS[5] (%)
No	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
(Fully) supervised	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
Self-supervised	De-noising (Vincent et al., 2010)	95.92±1.83	73.99±0.62	85.14±3.02	84.36±0.96	57.83±1.57
	In-painting (Pathak et al., 2016)	91.46±2.97	76.02±0.55	79.79±3.55	81.36±4.83	61.38±3.84
	Jigsaw (Noroozi and Favaro, 2016)	95.47±1.24	70.90±1.55	81.79±1.04	82.04±1.26	63.33±1.11
	DeepCluster (Caron et al., 2018)	97.22±0.55	74.95±0.46	84.82±0.62	82.66±1.00	65.96±0.85
	Patch shuffling (Chen et al., 2019a)	91.93±2.32	75.74±0.51	82.15±3.30	82.82±2.35	52.95±6.92
	Rubiks Cube (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

Genesis Chest CT

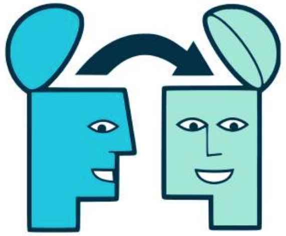

Target models

1. Zhou, Zongwei, et al. "Models genesis: Generic autodidactic models for 3d medical image analysis." International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, Cham, 2019.
2. Zhou, Zongwei, et al. "Models Genesis." arXiv preprint arXiv:2004.07882 (2020).

Aim #3: Extracting generic knowledge directly from unannotated images

Contribution: Models Genesis reduce annotation efforts by at least 30%

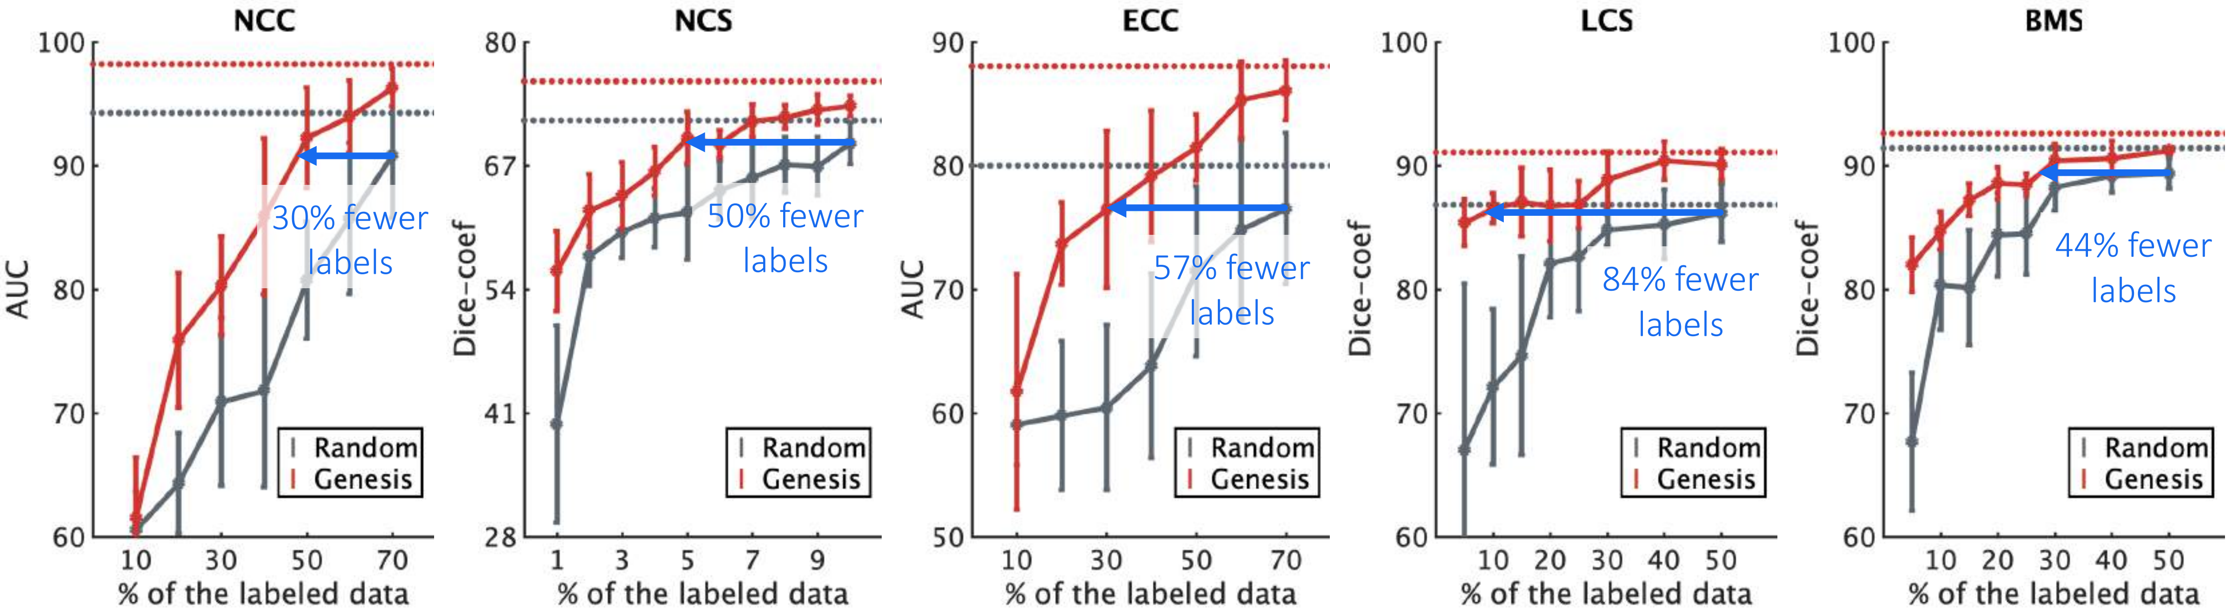

[1]NCC Lung nodule false positive reduction in CT images
[2]NCS Lung nodule segmentation in CT images
[3]ECC Pulmonary embolism false positive reduction in CT images
[4]LCS Liver segmentation in CT images
[5]BMS Brain tumor segmentation in MR images

Genesis Chest CT 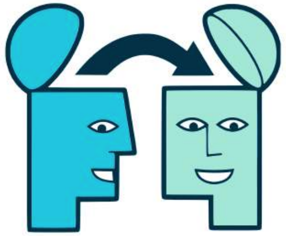 Target models

1. Zhou, Zongwei, et al. "Models genesis: Generic autodidactic models for 3d medical image analysis." International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, Cham, 2019.
2. Zhou, Zongwei, et al. "Models Genesis." arXiv preprint arXiv:2004.07882 (2020).

Aim #3: Extracting generic knowledge directly from unannotated images

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

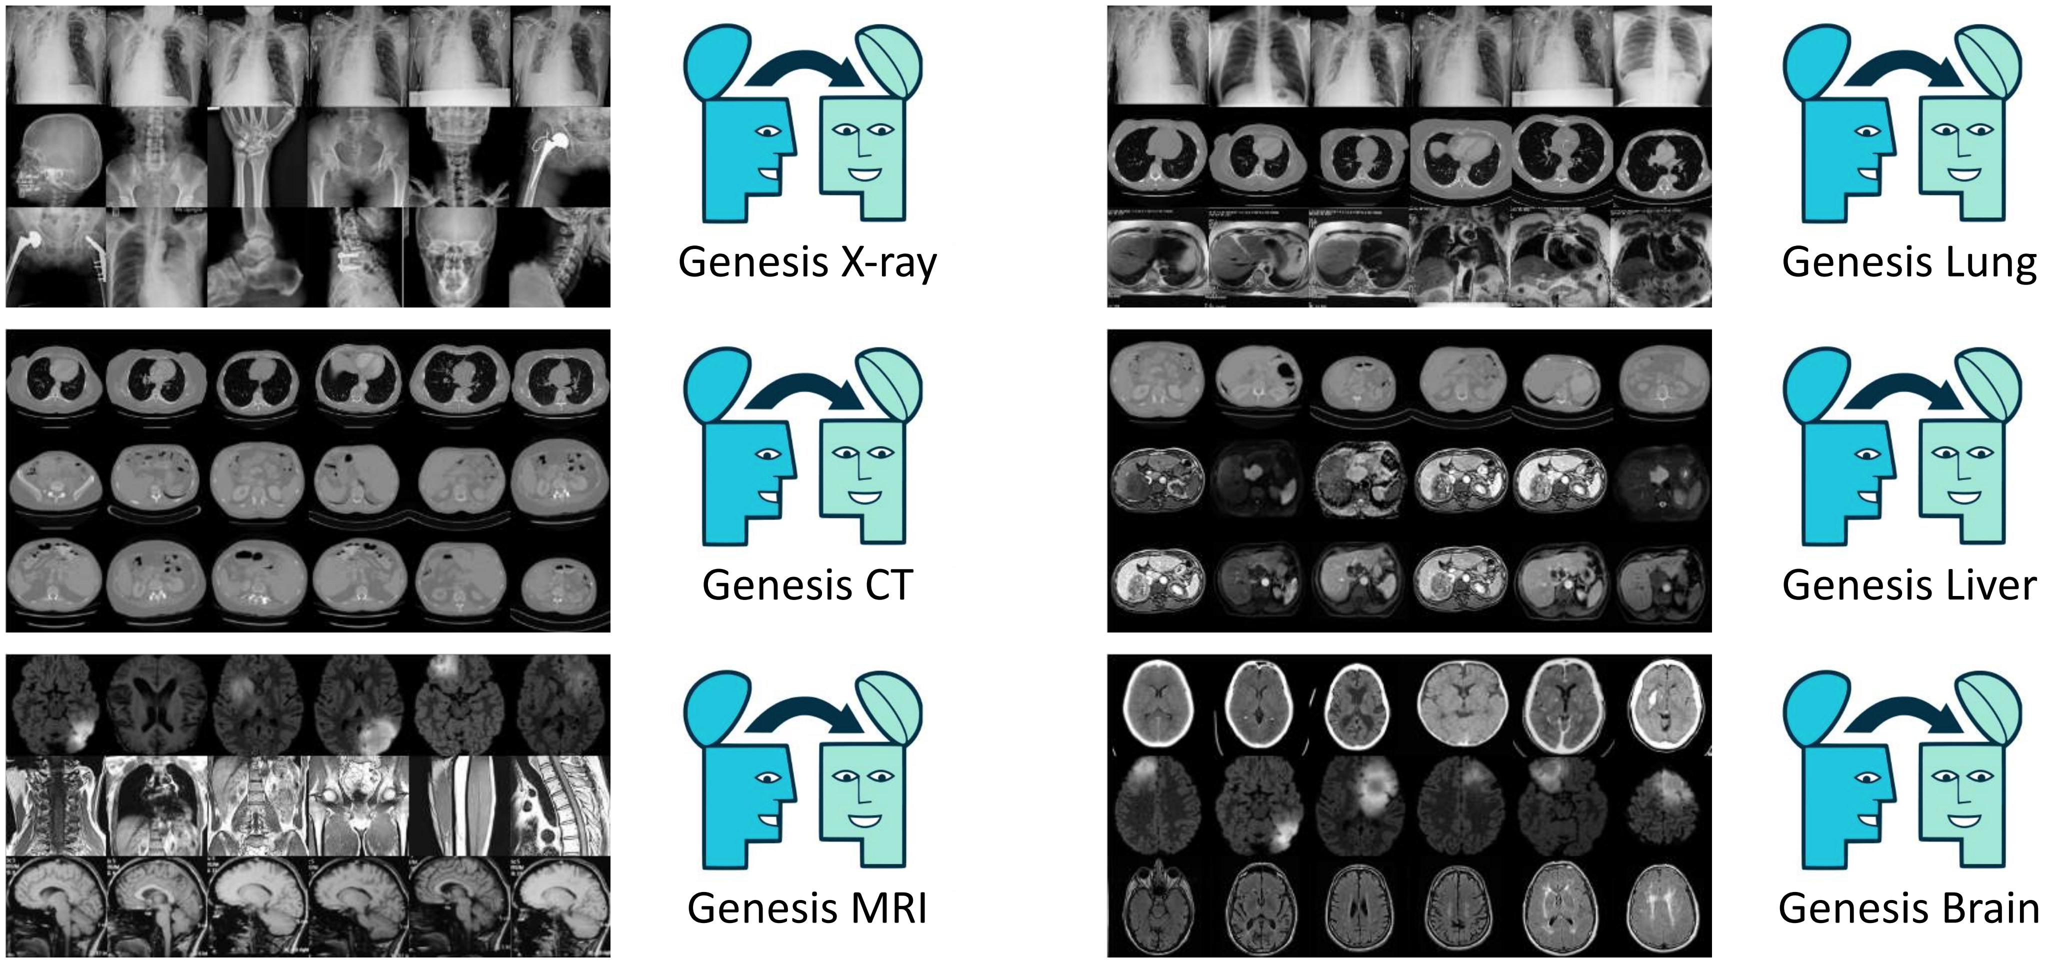

Holy Grail: effective across diseases, organs, and modalities.

Featured Publications for Aim #3:

1. Z. Zhou, V. Sodha, M. M. Rahman Siddiquee, R. Feng, N. Tajbakhsh, M. Gotway, J. Liang, 2019. Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis. MICCAI'19, Young Scientist Award, Best Presentation Award Finalist.
2. Z. Zhou, V. Sodha, J. Pang, M. Gotway, J. Liang, 2020. Models Genesis. Medical Image Analysis, MedIA Best Paper Award.

Holy Grail: effective across diseases, organs, and modalities.

Clinical Impacts of Aim #3:

1. Transfer learning can greatly reduce the cost and effort required to build a dataset and retrain the model. Instead of building a model from scratch (demanding numerous data acquisition and annotation), a smaller dataset can be used to efficiently fine-tune the existing model.
2. Generic pre-trained models can serve as a primary source of transfer learning for many medical imaging applications, leading to accelerated training and improved performance.

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

Aim #1: Acquiring necessary annotation efficiently from human experts

Aim #2: Utilizing existing annotation effectively from advanced architecture

Aim #3: Extracting generic knowledge directly from unannotated images

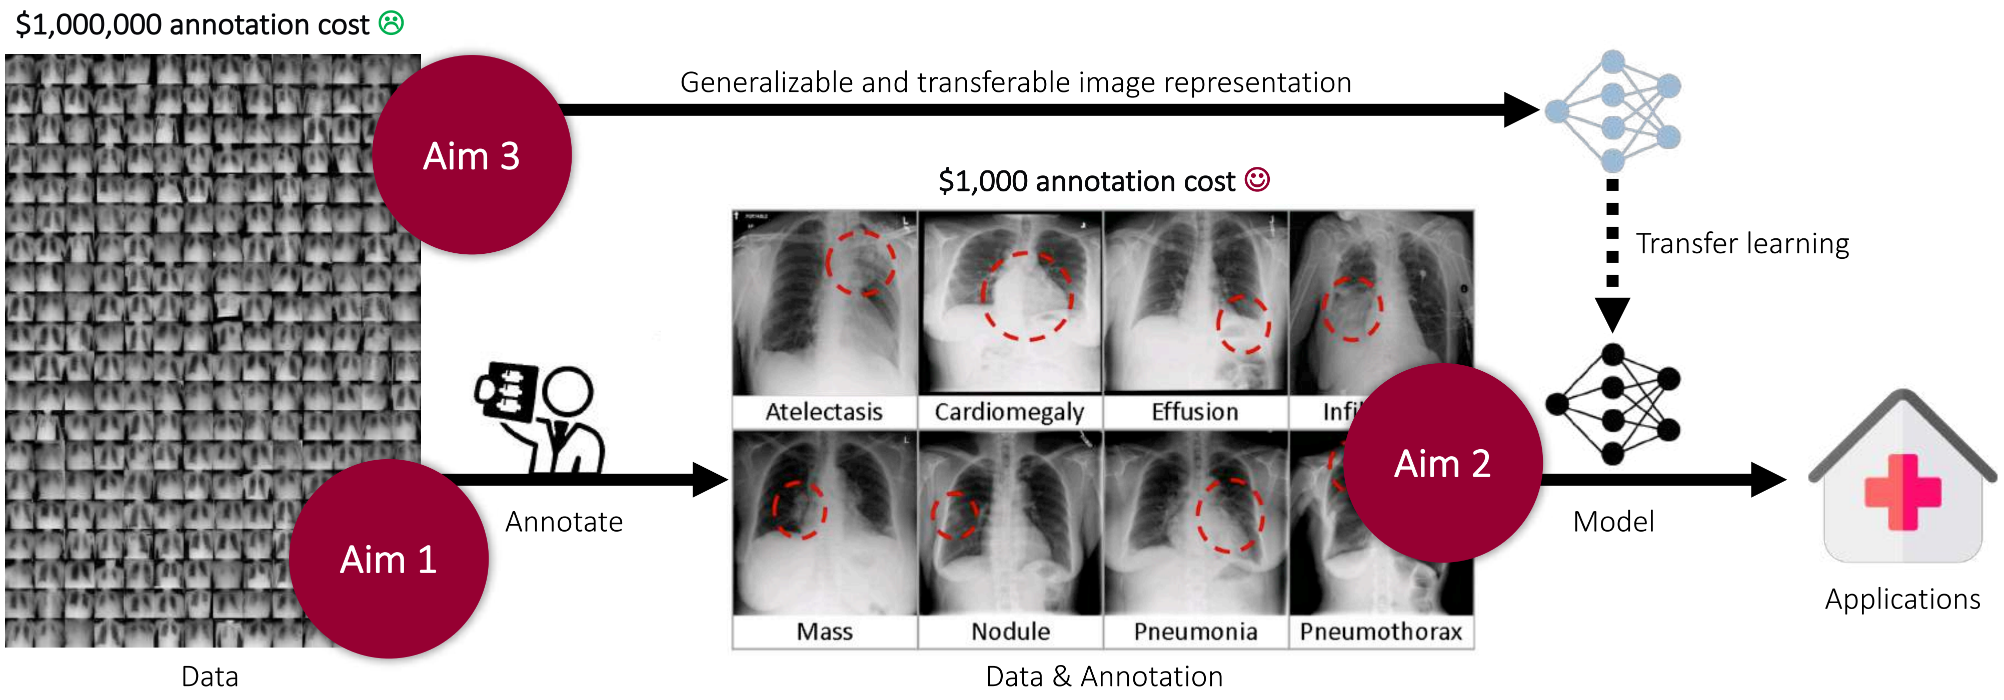

Acknowledgements

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

Funding for research program supported by

- NIH R01 (R01HL128785)
- ASU-Mayo Grant
- Mayo Innovation Grant

Annotation-efficient Deep Learning for Computer-aided Diagnosis in Medical Imaging

Zongwei Zhou

Ph.D. Candidate, Biomedical Informatics

College of Health Solutions, Arizona State University

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