We provide pre-trained 3D models!

Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis

MICCAI 2019 Young Scientist Award

Zongwei Zhou[1], Vatsal Sodha[1], Md Mahfuzur Rahman Siddiquee[1], Ruibin Feng[1], Nima Tajbakhsh[1], Michael B. Gotway[2], and Jianming Liang[1]

[1] Arizona State University [2] Mayo Clinic

We provide pre-trained 3D models!

Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis

MICCAI 2019 Young Scientist Award

“这篇文章的贡献是设计了一个针对三维医学图像分析的预训练模型，
这样解决了以前大家只能用 ImageNet 里的二维数据训练出来的预模型，
并且得到更好的效果；在5个医学图像的分割和分类问题上取得领先的效果；
在作者的口头发言中也给出了开源代码
（https://github.com/MrGiovanni/ModelsGenesis）。”

——沈定刚教授

MICCAI 2019：纪录、风向与学术思考

Existing publicly available 3D pre-trained models?

deepmind / **kinetics-i3d** 240K annotated videos

Watch 52 | Unstar 1.1k | Fork 355

<> Code | Issues 68 | Pull requests 4 | Actions | Projects 0 | Wiki | Security | Insights

Convolutional neural network model for video classification trained on the Kinetics dataset.

NifTK / **NiftyNet** 90 annotated subjects

Used by 19 | Watch 89 | Unstar 1.1k | Fork 365

<> Code | Issues 95 | Pull requests 3 | Actions | Projects 3 | Wiki | Security | Insights

An open-source convolutional neural networks platform for research in medical image analysis and image-guided therapy
http://niftynet.io

Tencent / **MedicalNet** 1,638 annotated subjects

Unwatch 43 | Unstar 790 | Fork 210

<> Code | Issues 28 | Pull requests 2 | Actions | Projects 0 | Wiki | Security | Insights

Many studies have shown that the performance on deep learning is significantly affected by volume of training data. The MedicalNet project provides a series of 3D-ResNet pre-trained models and relative code.

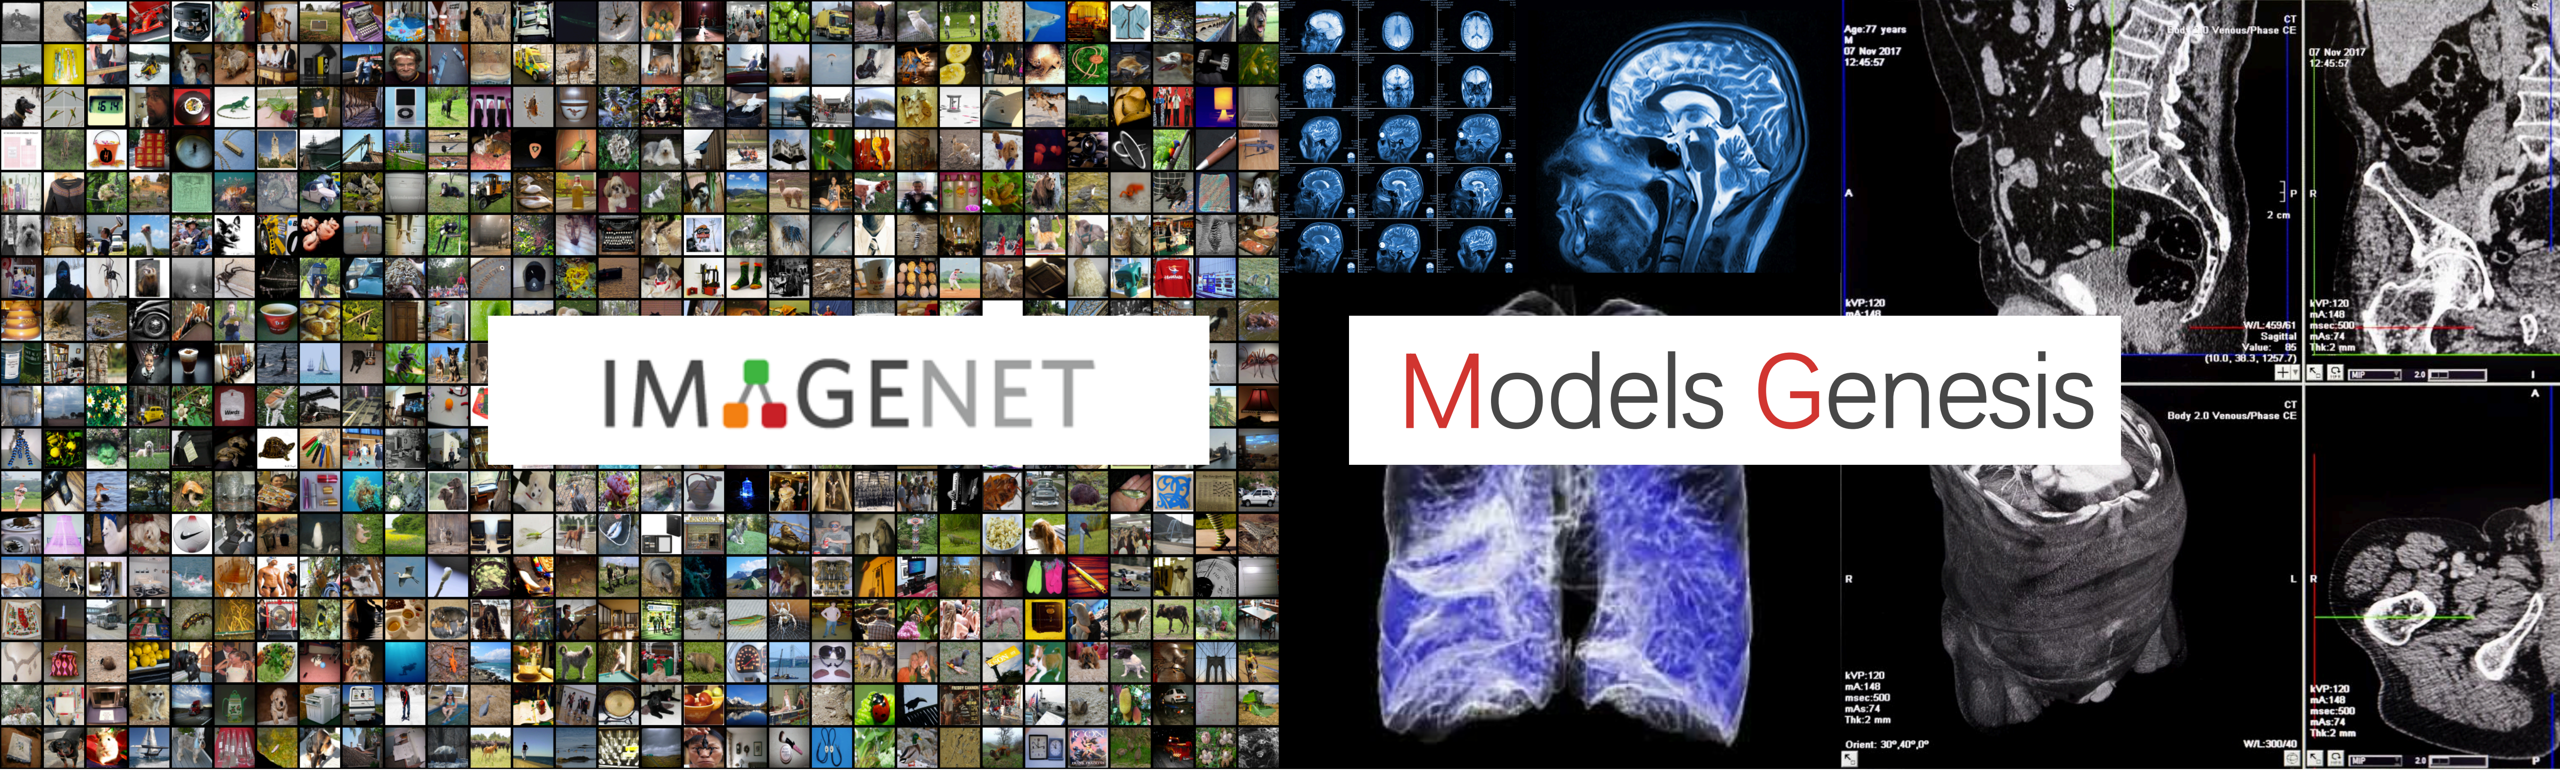

Natural images	Medical images
Formed in 2D	Formed in 3D
>14,000,000 annotation	Zero annotation

ImageNet demands huge amount of annotation efforts,
but Models Genesis are pre-trained with self-supervision.

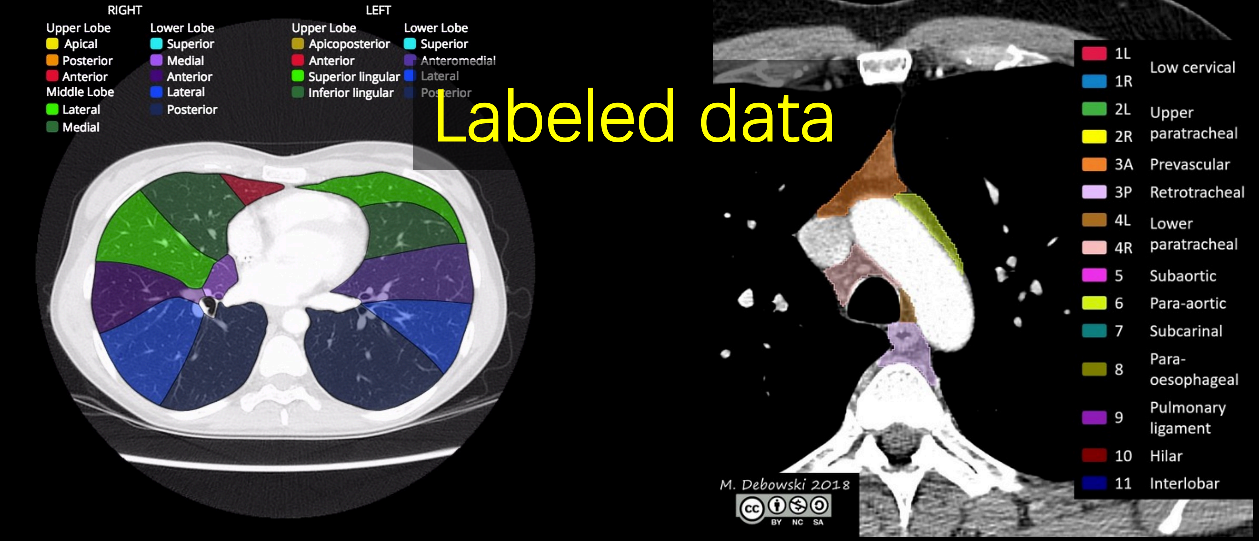

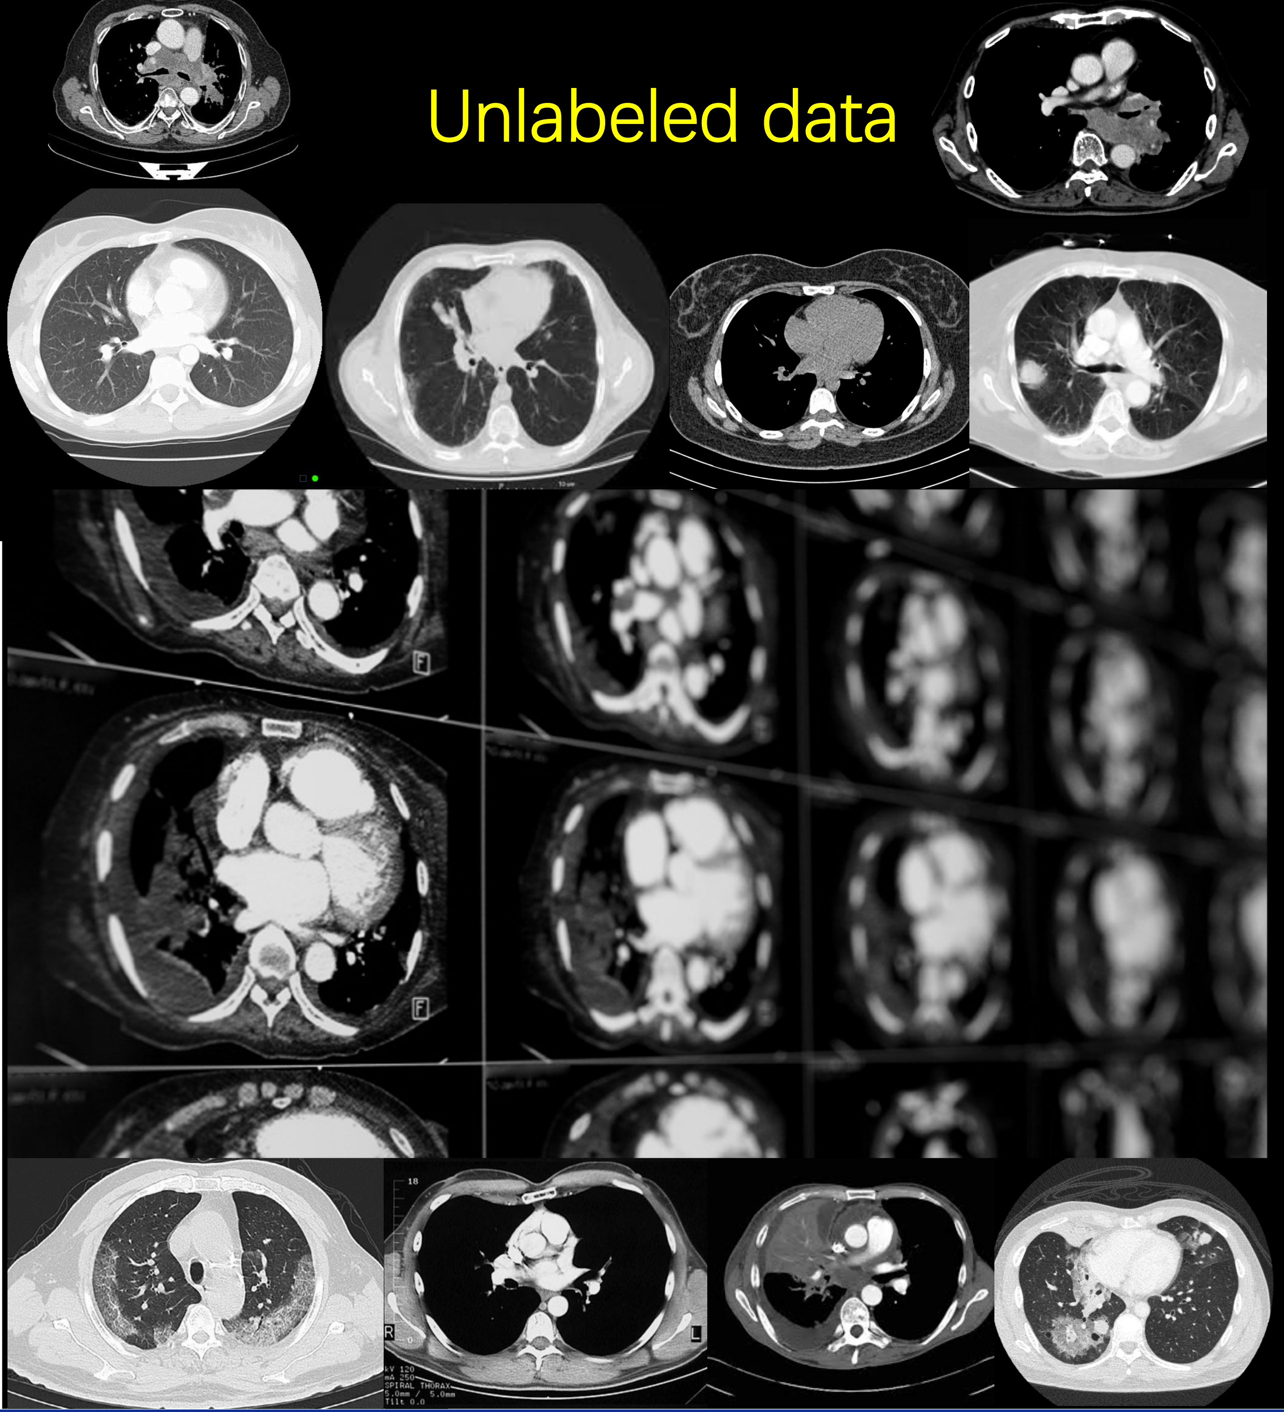

Can we utilize the large number of available Chest CT images <u>without systematic annotation</u> to train source models that can yield high-performance target models via transfer learning?

We design it as a simple image restoration task, through which, the model can learn representation directly from image data itself.

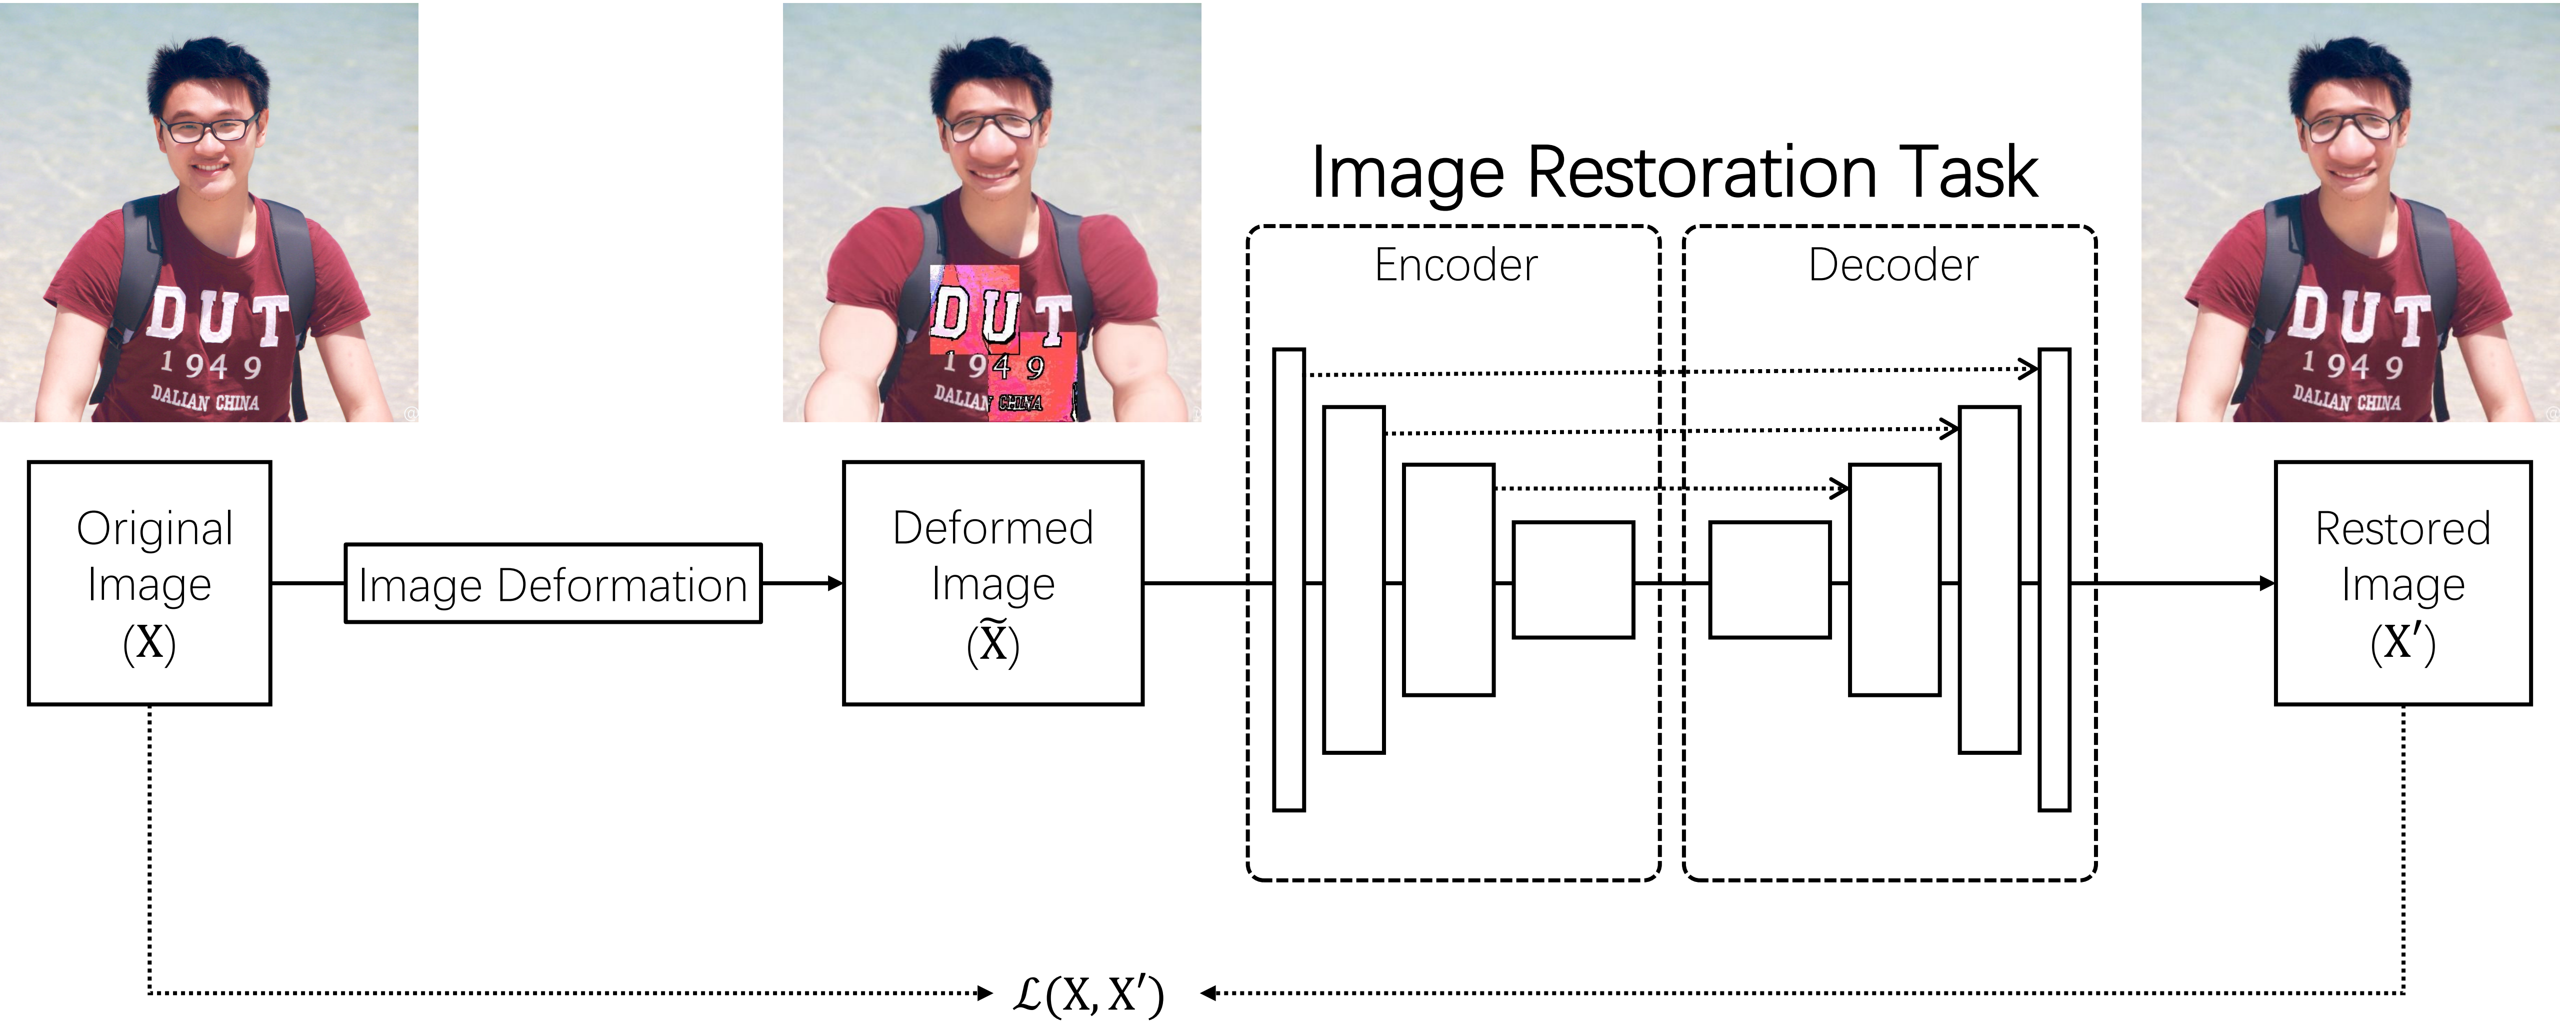

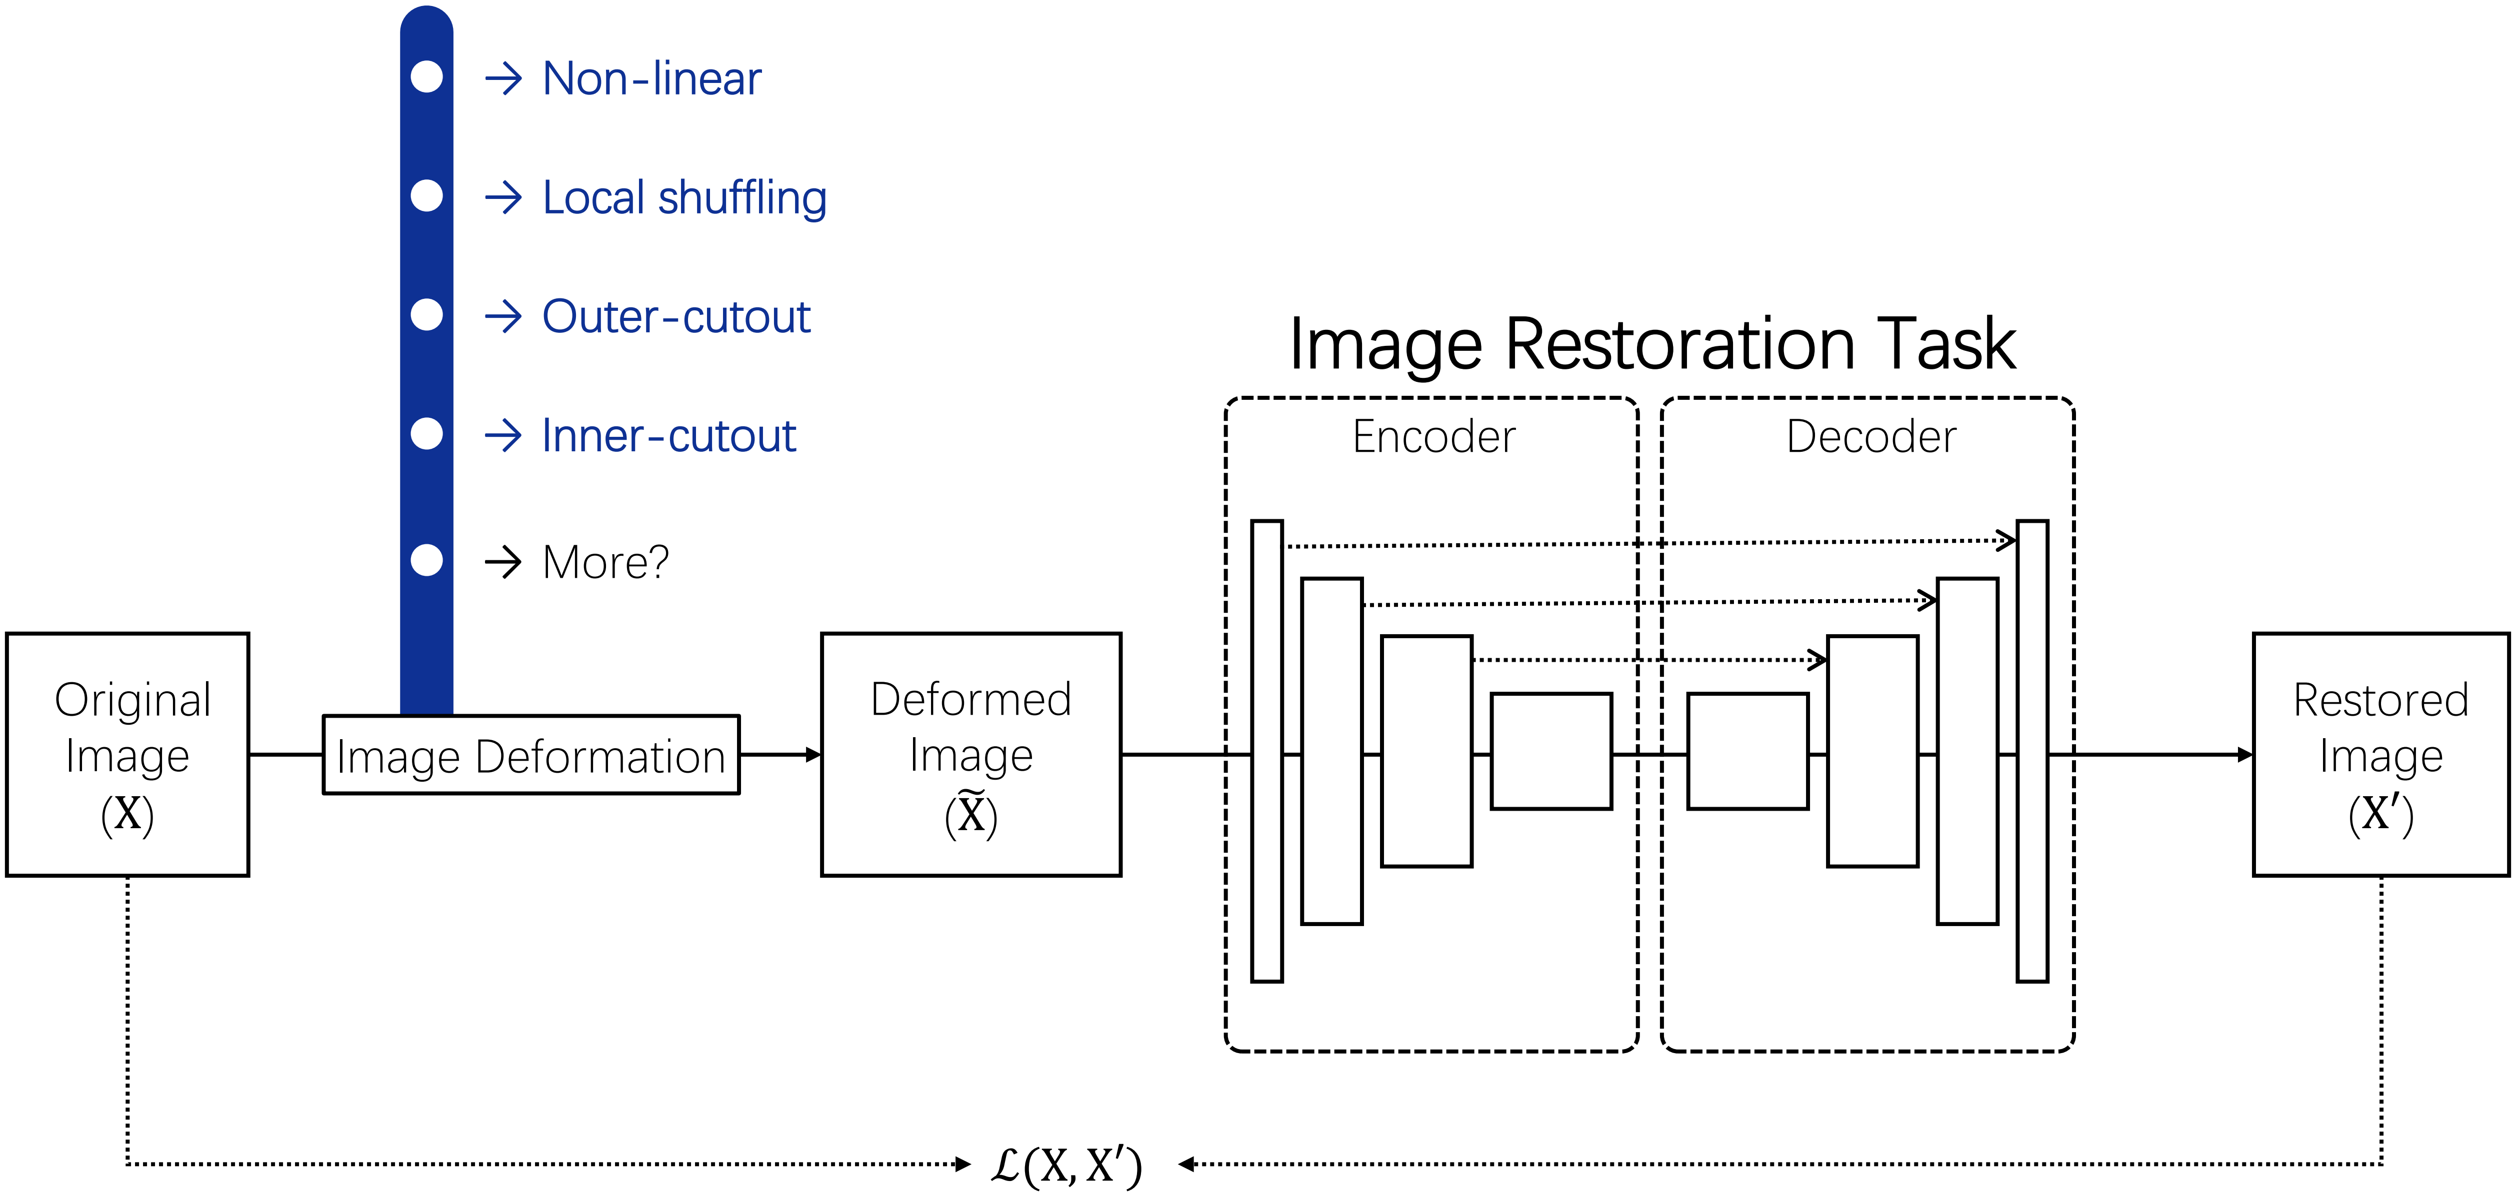
Non-linear
Local shuffling
Outer-cutout
Inner-cutout
More?
Image Restoration Task
Encoder
Decoder
Original Image (X)
Image Deformation
Deformed Image ($\tilde{X}$)
Restored Image (X′)
$\mathcal{L}(X, X')$

I. Non-linear transformation

CT scan itself naturally comes with the *pixel-wise* annotation

Substance		Hounsfield units (HU)
Air		-1000
Fat		-120 to -90
Water		0
Bone	Cancellous	+300 to +400
	Cortical	+1800 to +1900
Parenchyma	Lung	-700 to −600
	Kidney	+20 to +45
	Liver	+54 to +66
	Lymph nodes	+10 to +20
	Muscle	+35 to +55

Source from en.wikipedia.org/wiki/Hounsfield_scale

I. Non-linear transformation

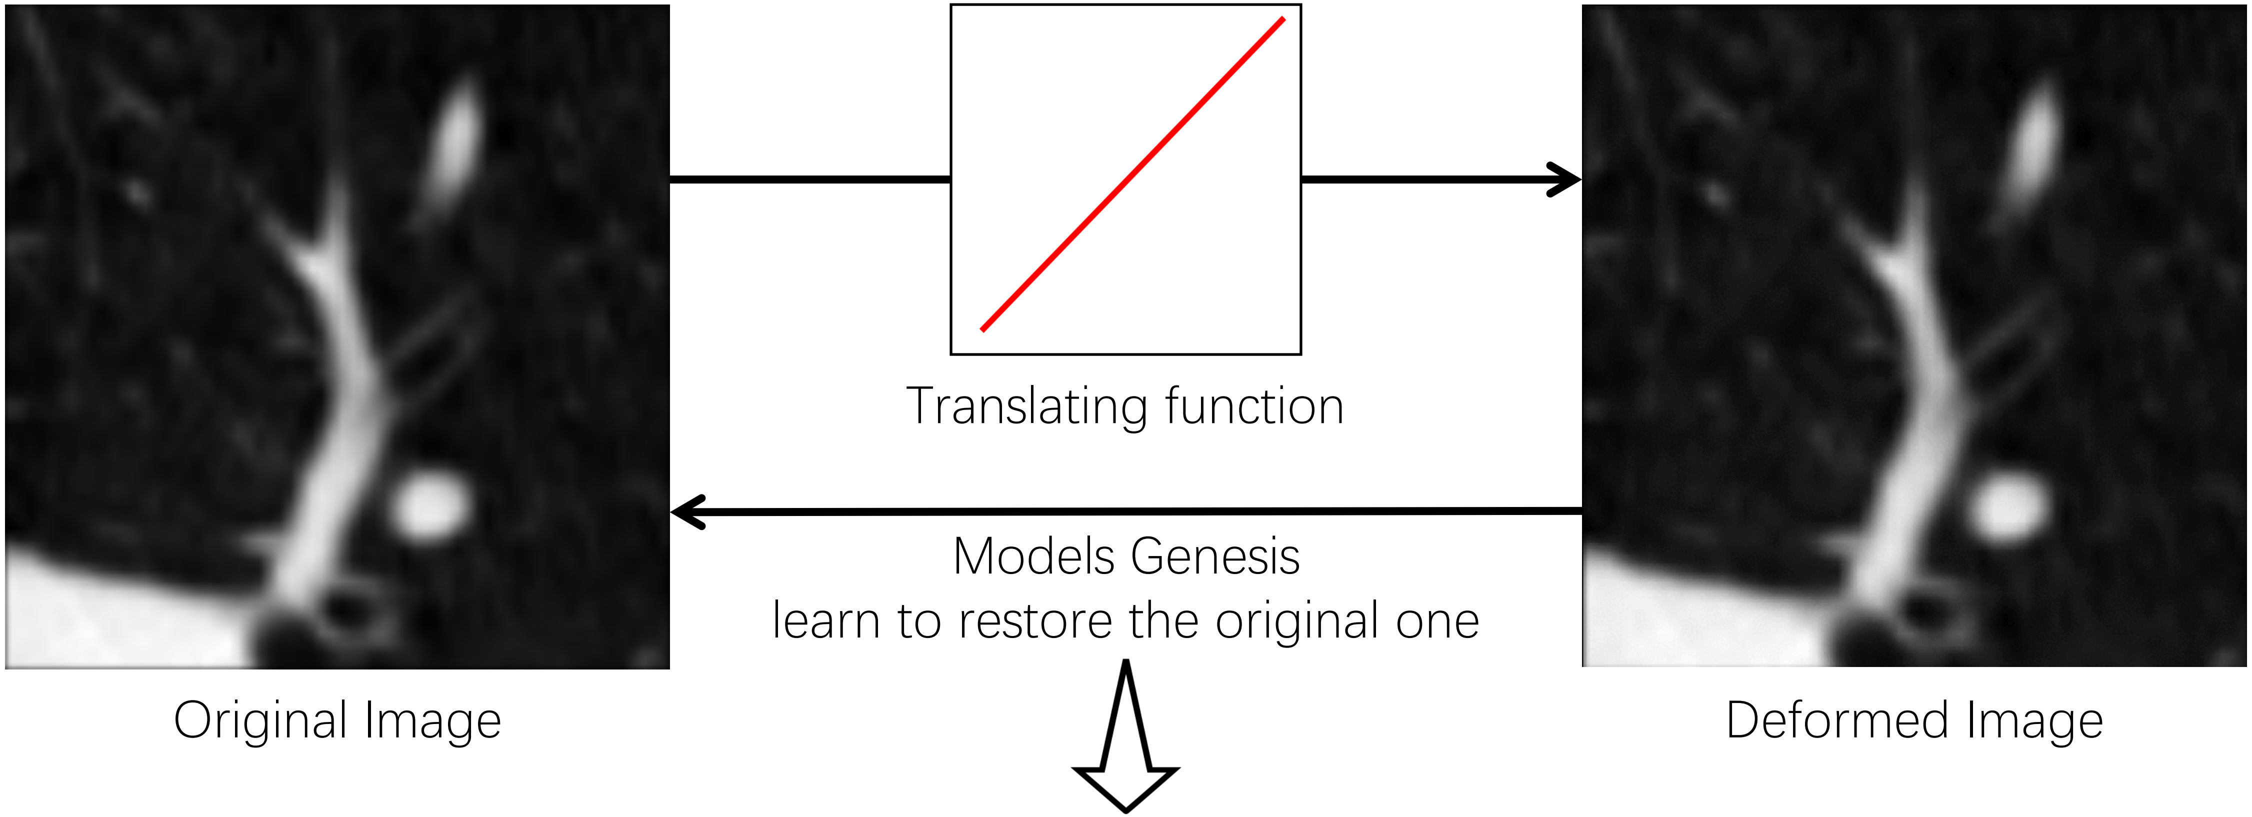

Learning organ appearance and intensity distribution

II. Local pixel shuffling

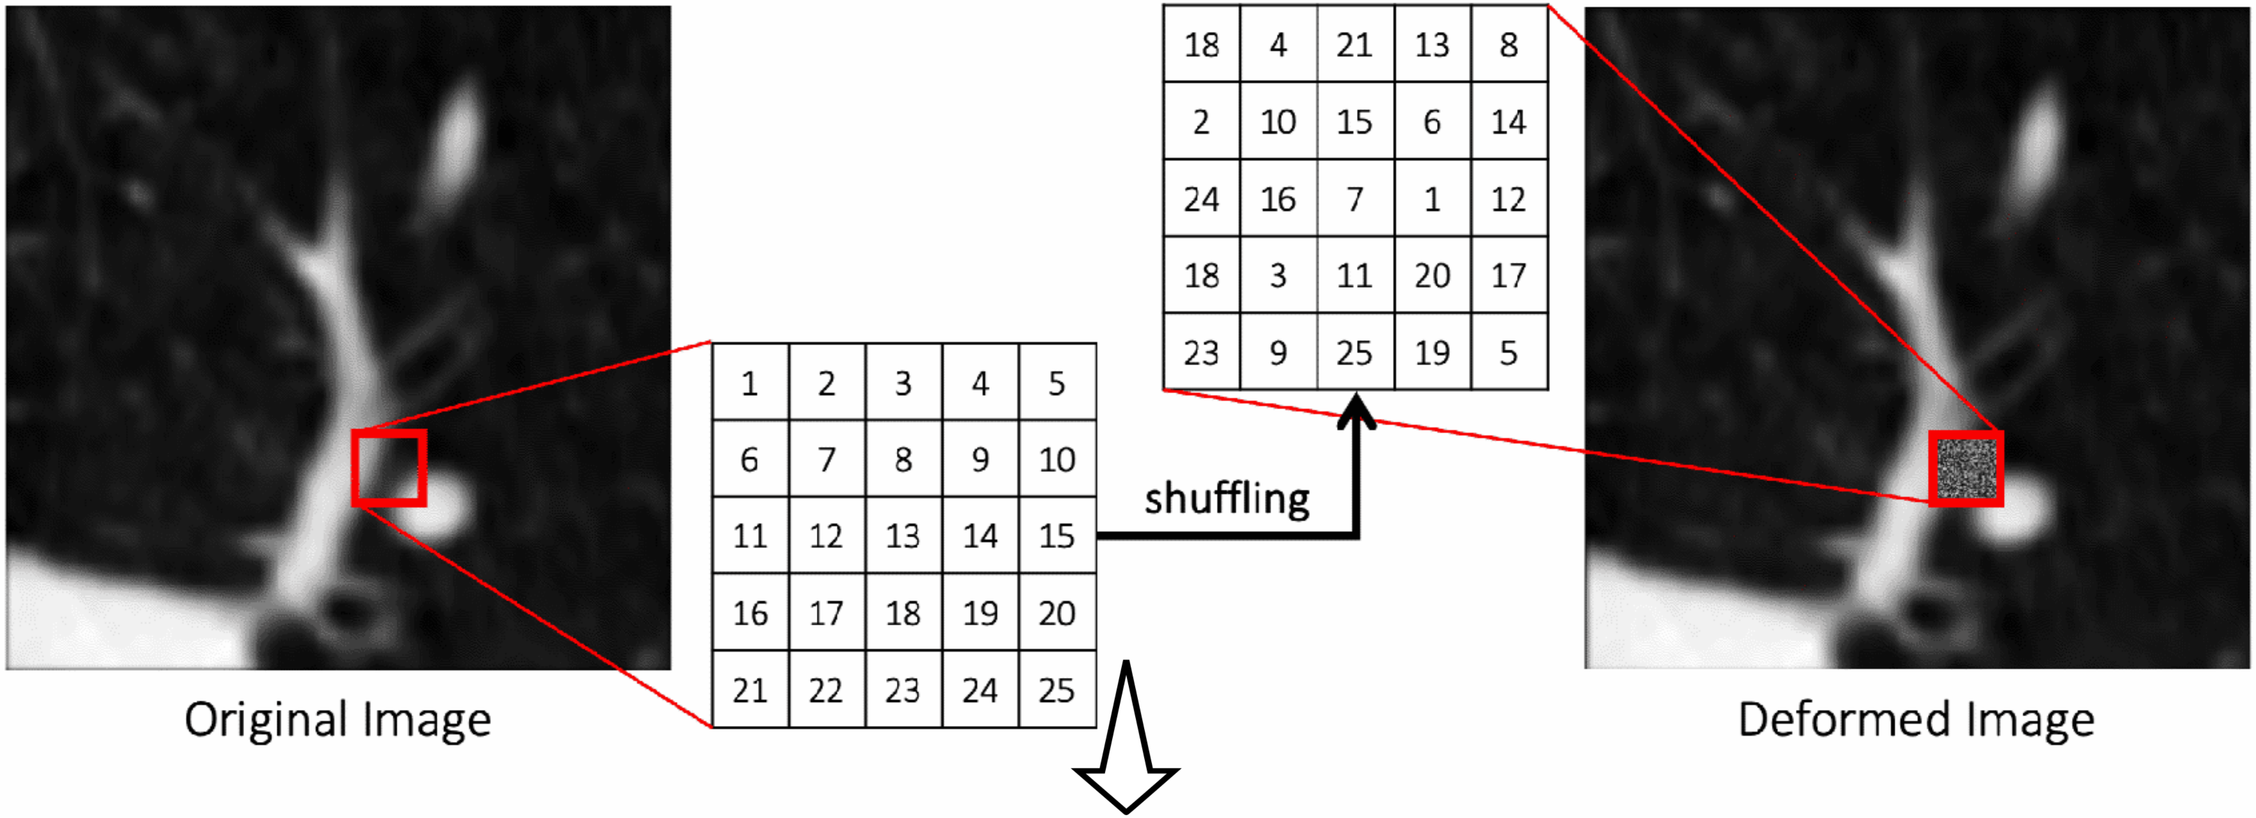

Original Image

Deformed Image

Learning organ texture and local boundaries

III. Outer-cutout

IV. Inner-cutout

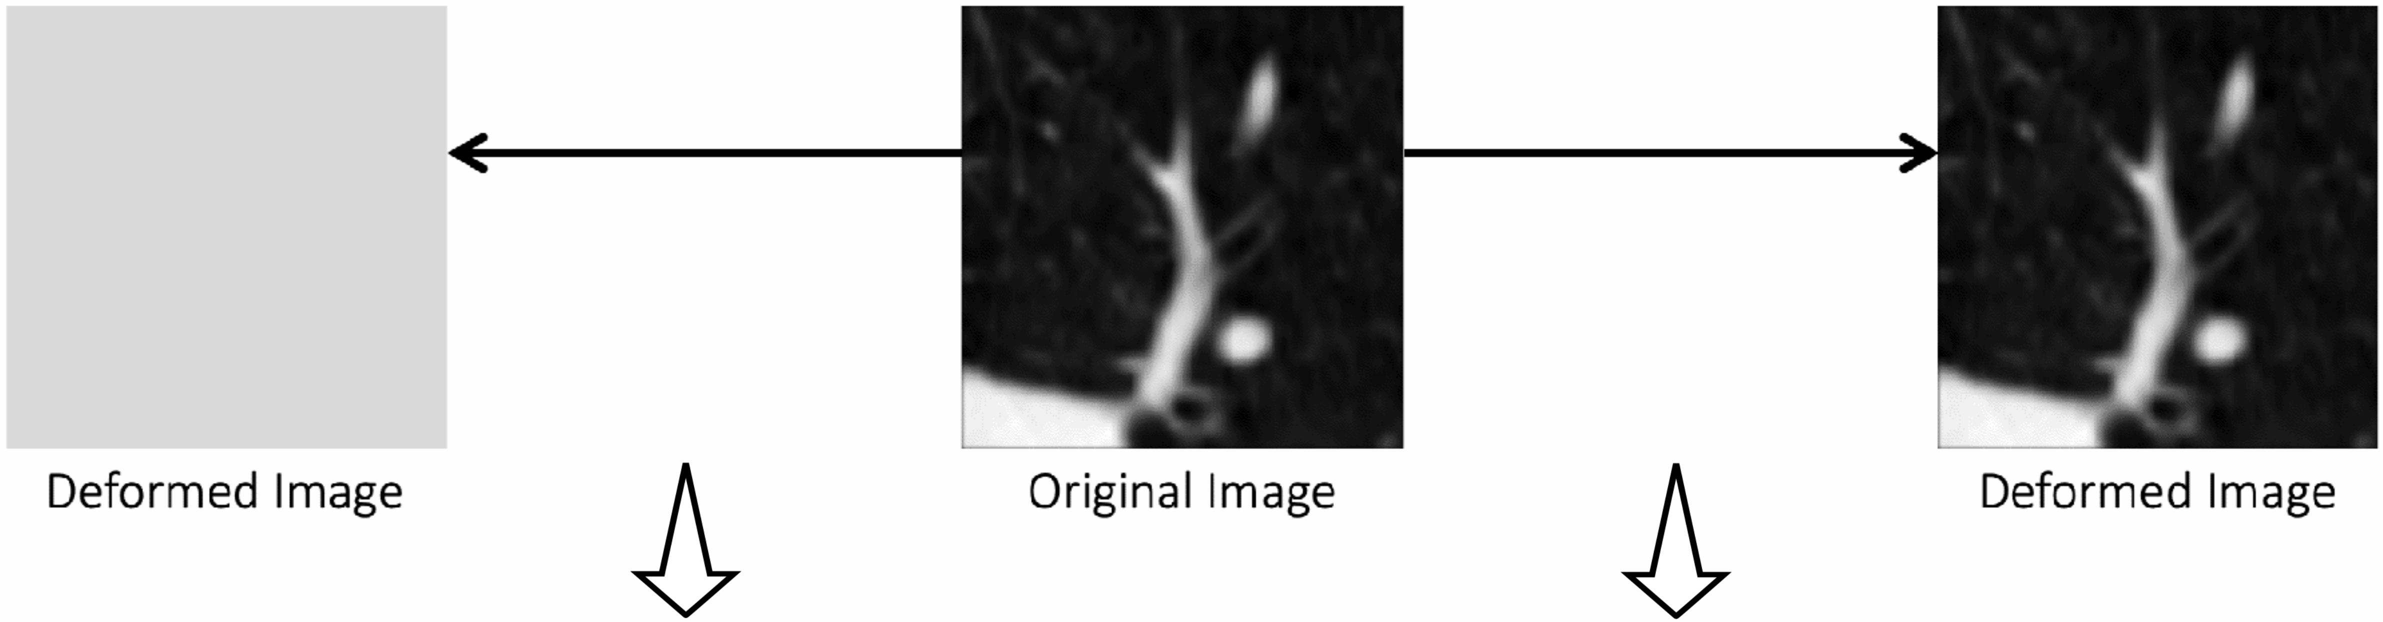

Learning organ spatial layout and global geometry

Learning local continuities

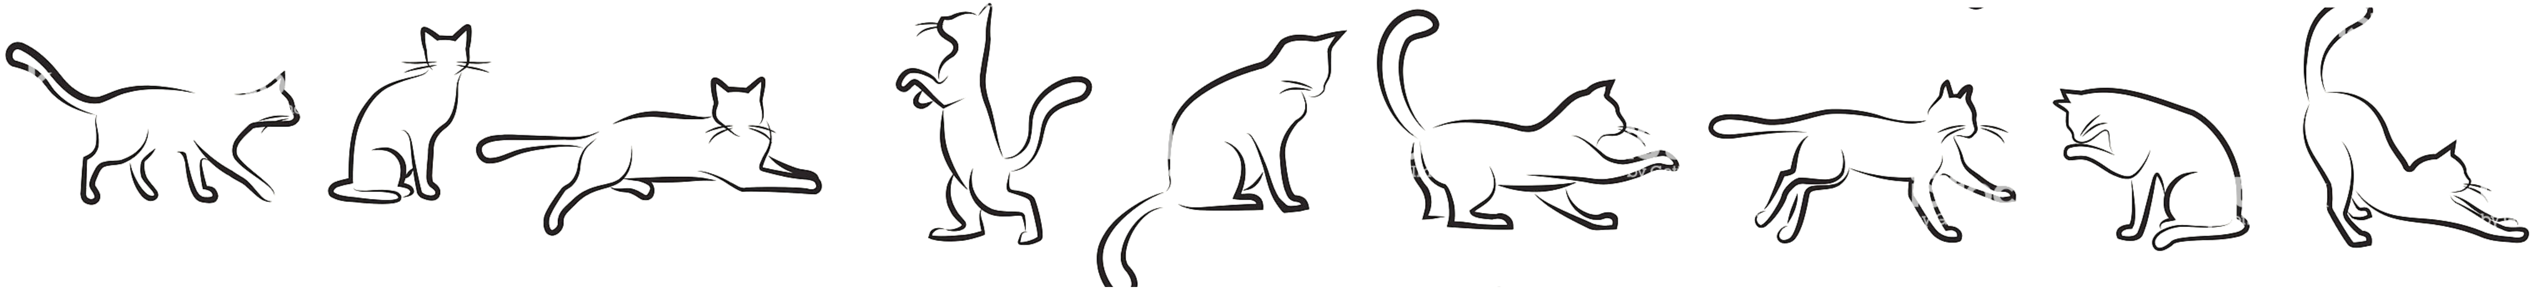

Why are the proposed image deformations in your paper effective?

Medical images contain <u>similar anatomy</u>. The sophisticated yet recurrent anatomy offers <u>consistent patterns</u> for self-supervised learning to discover common representation of a particular body part.

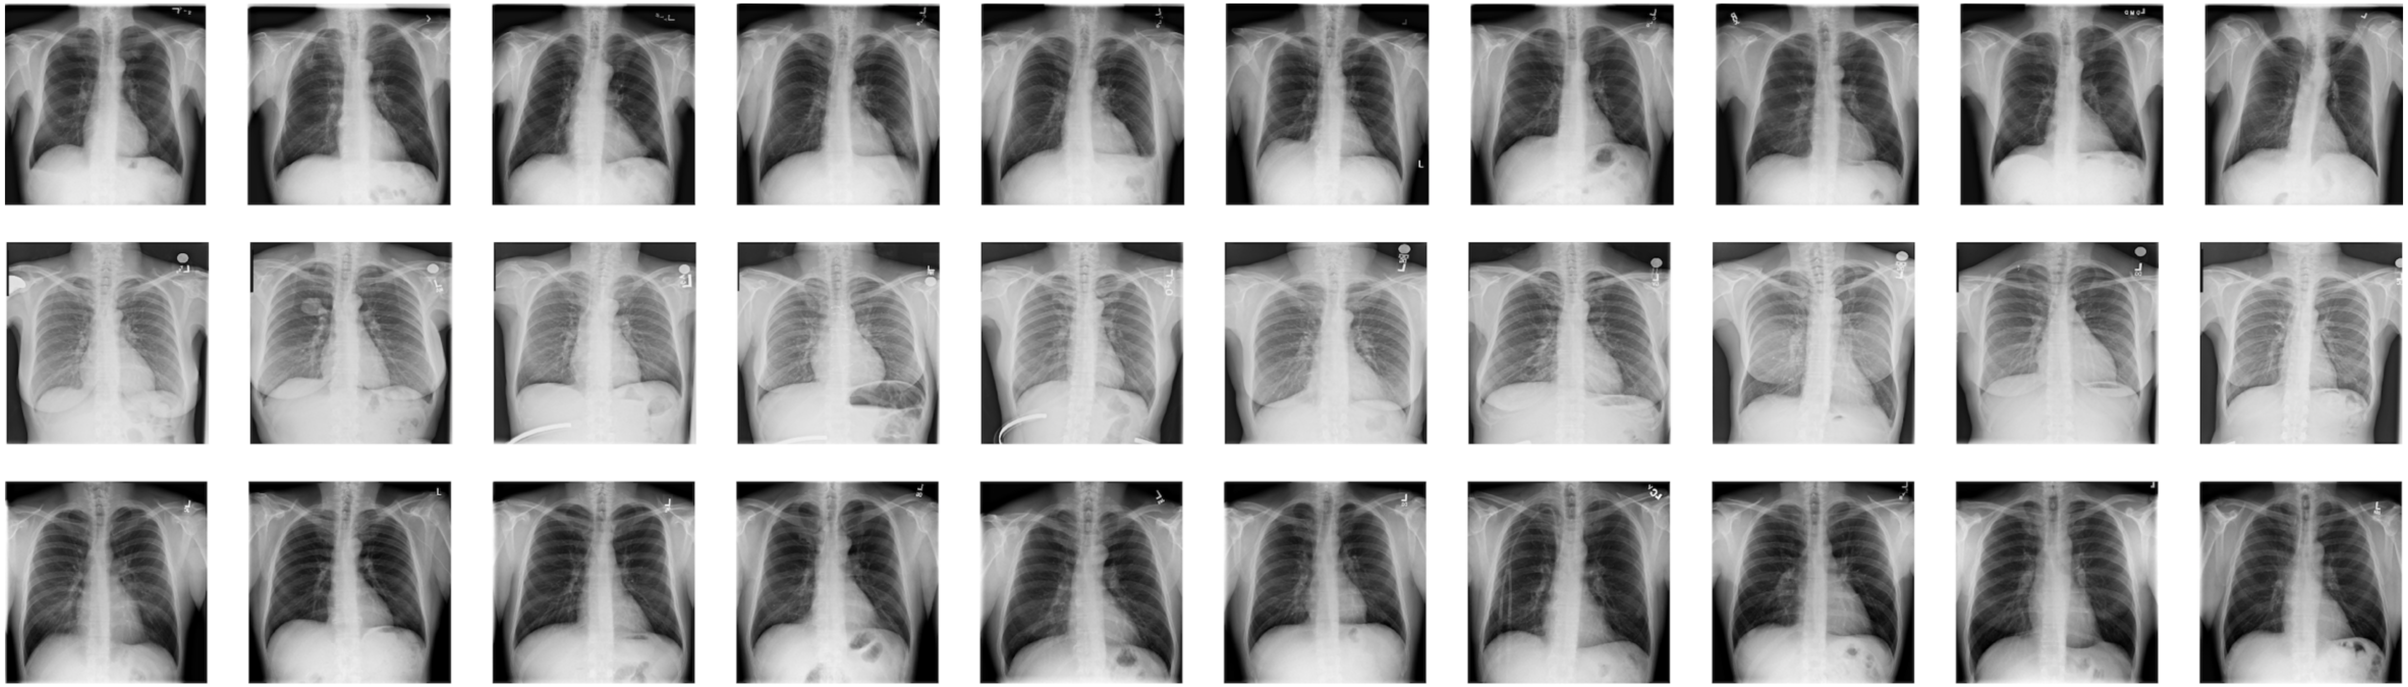

Once pre-trained,

the **encoder** could be used
for target classification tasks
e.g., brain tumor classification;

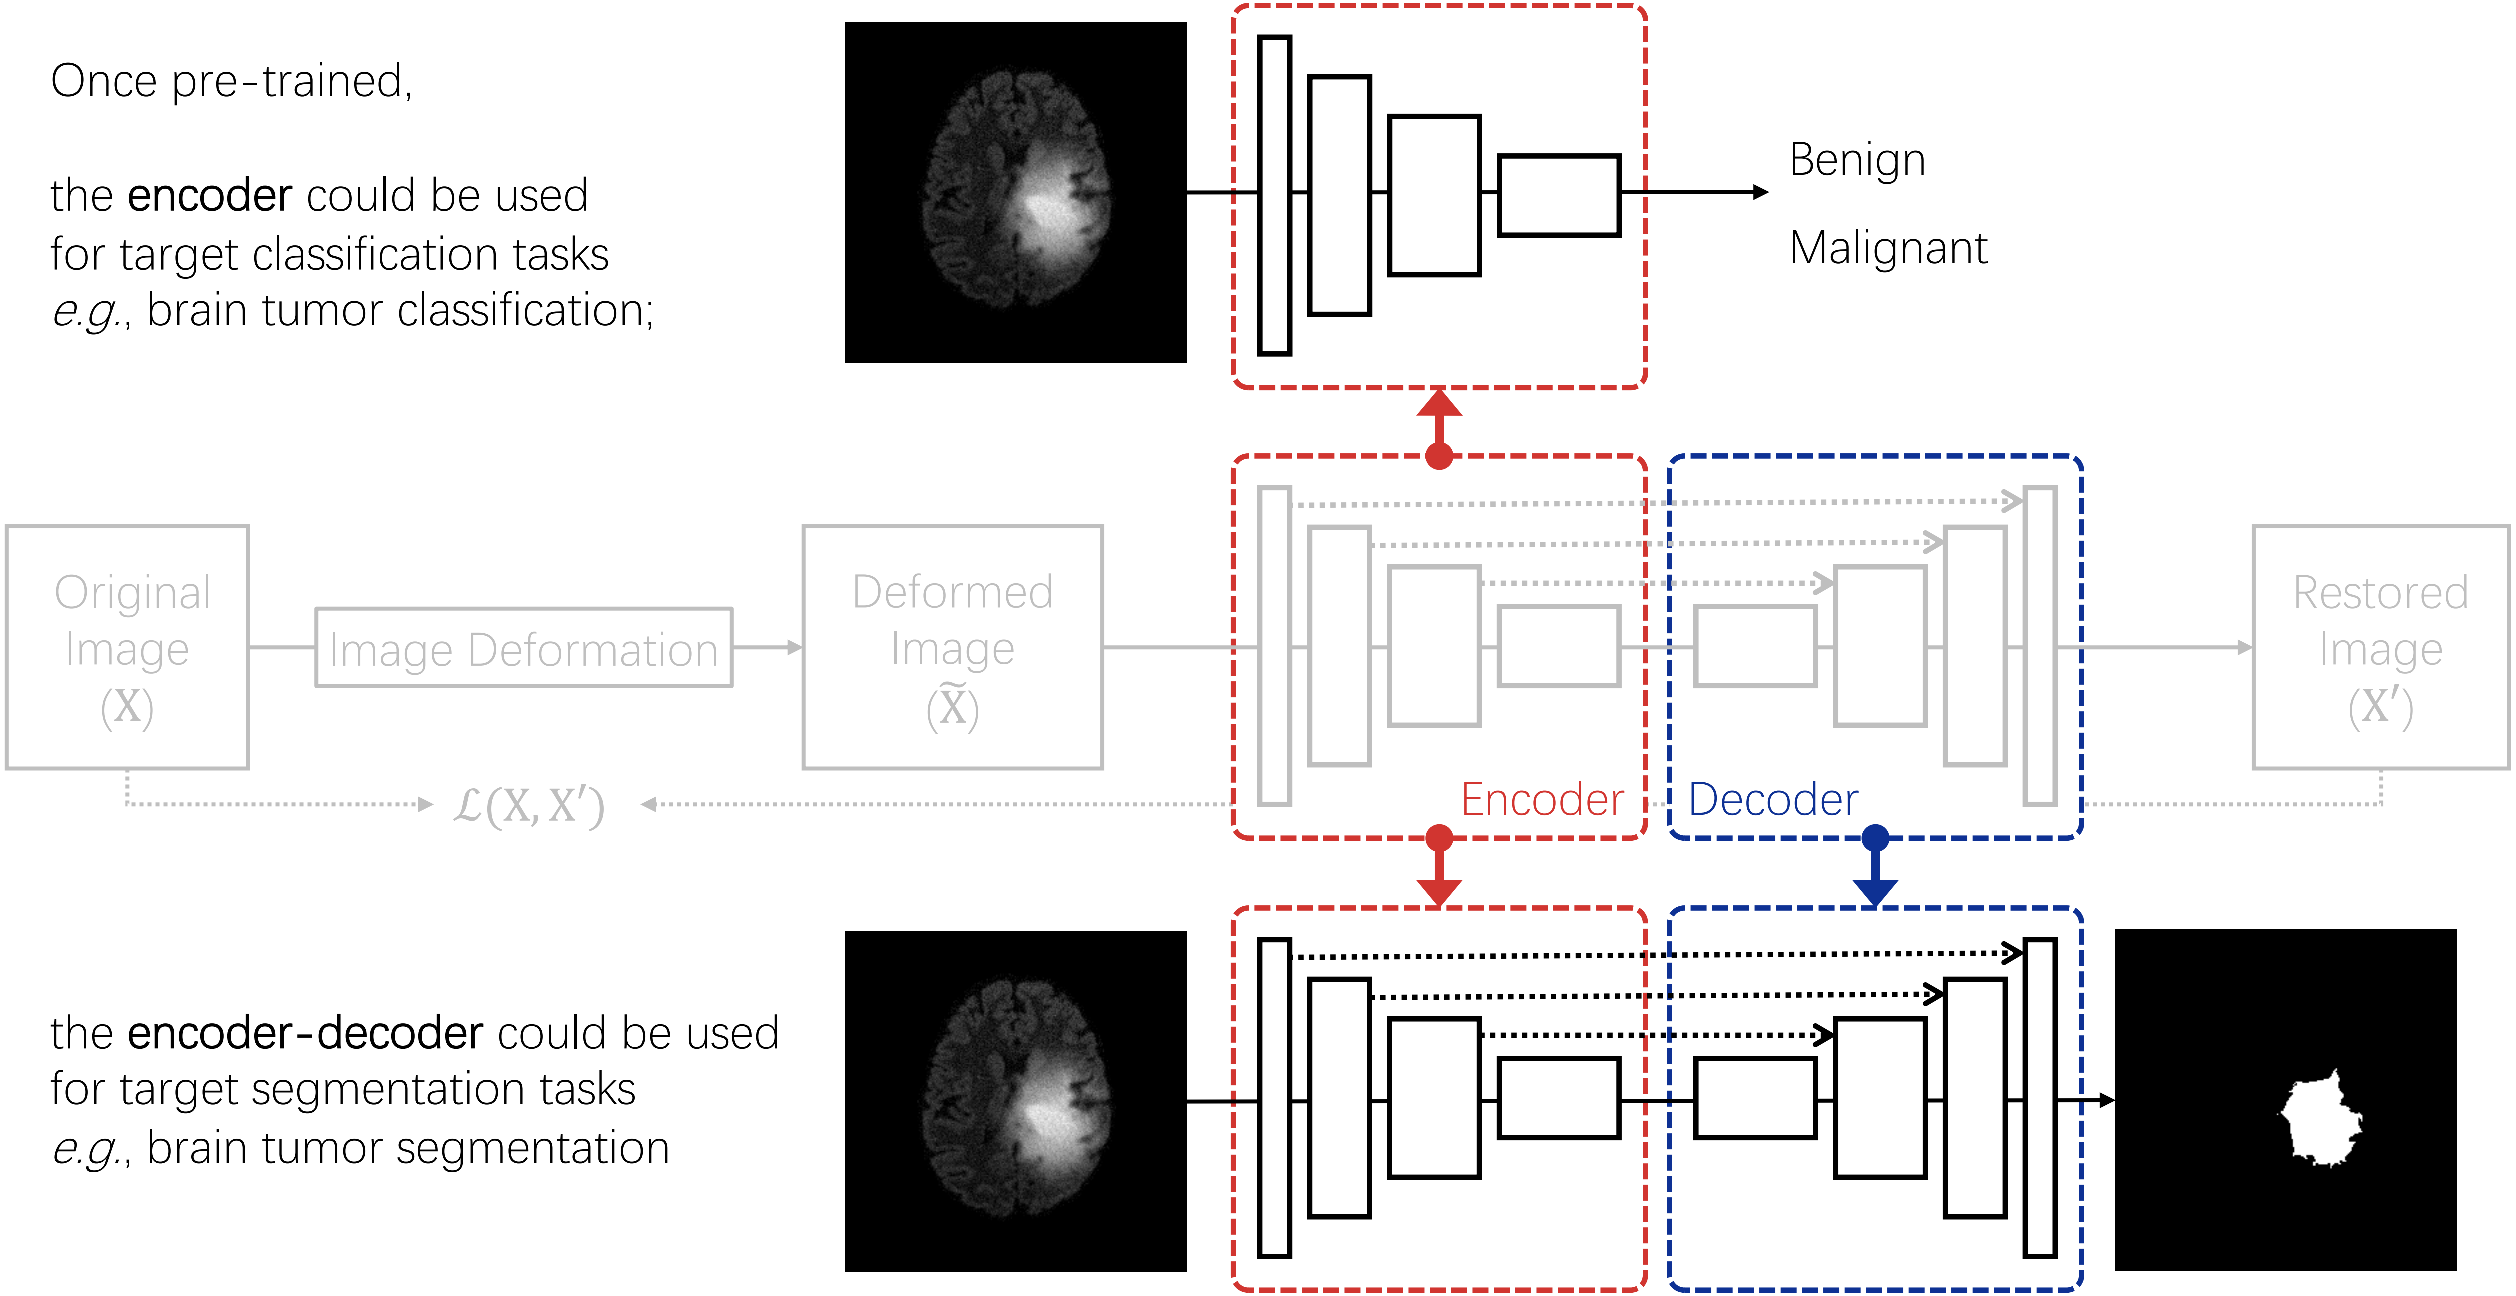

the **encoder-decoder** could be used
for target segmentation tasks
e.g., brain tumor segmentation

Two classification applications

1. Lung nodule false positive reduction (CT)
2. PE false positive reduction 3D CT (CT)

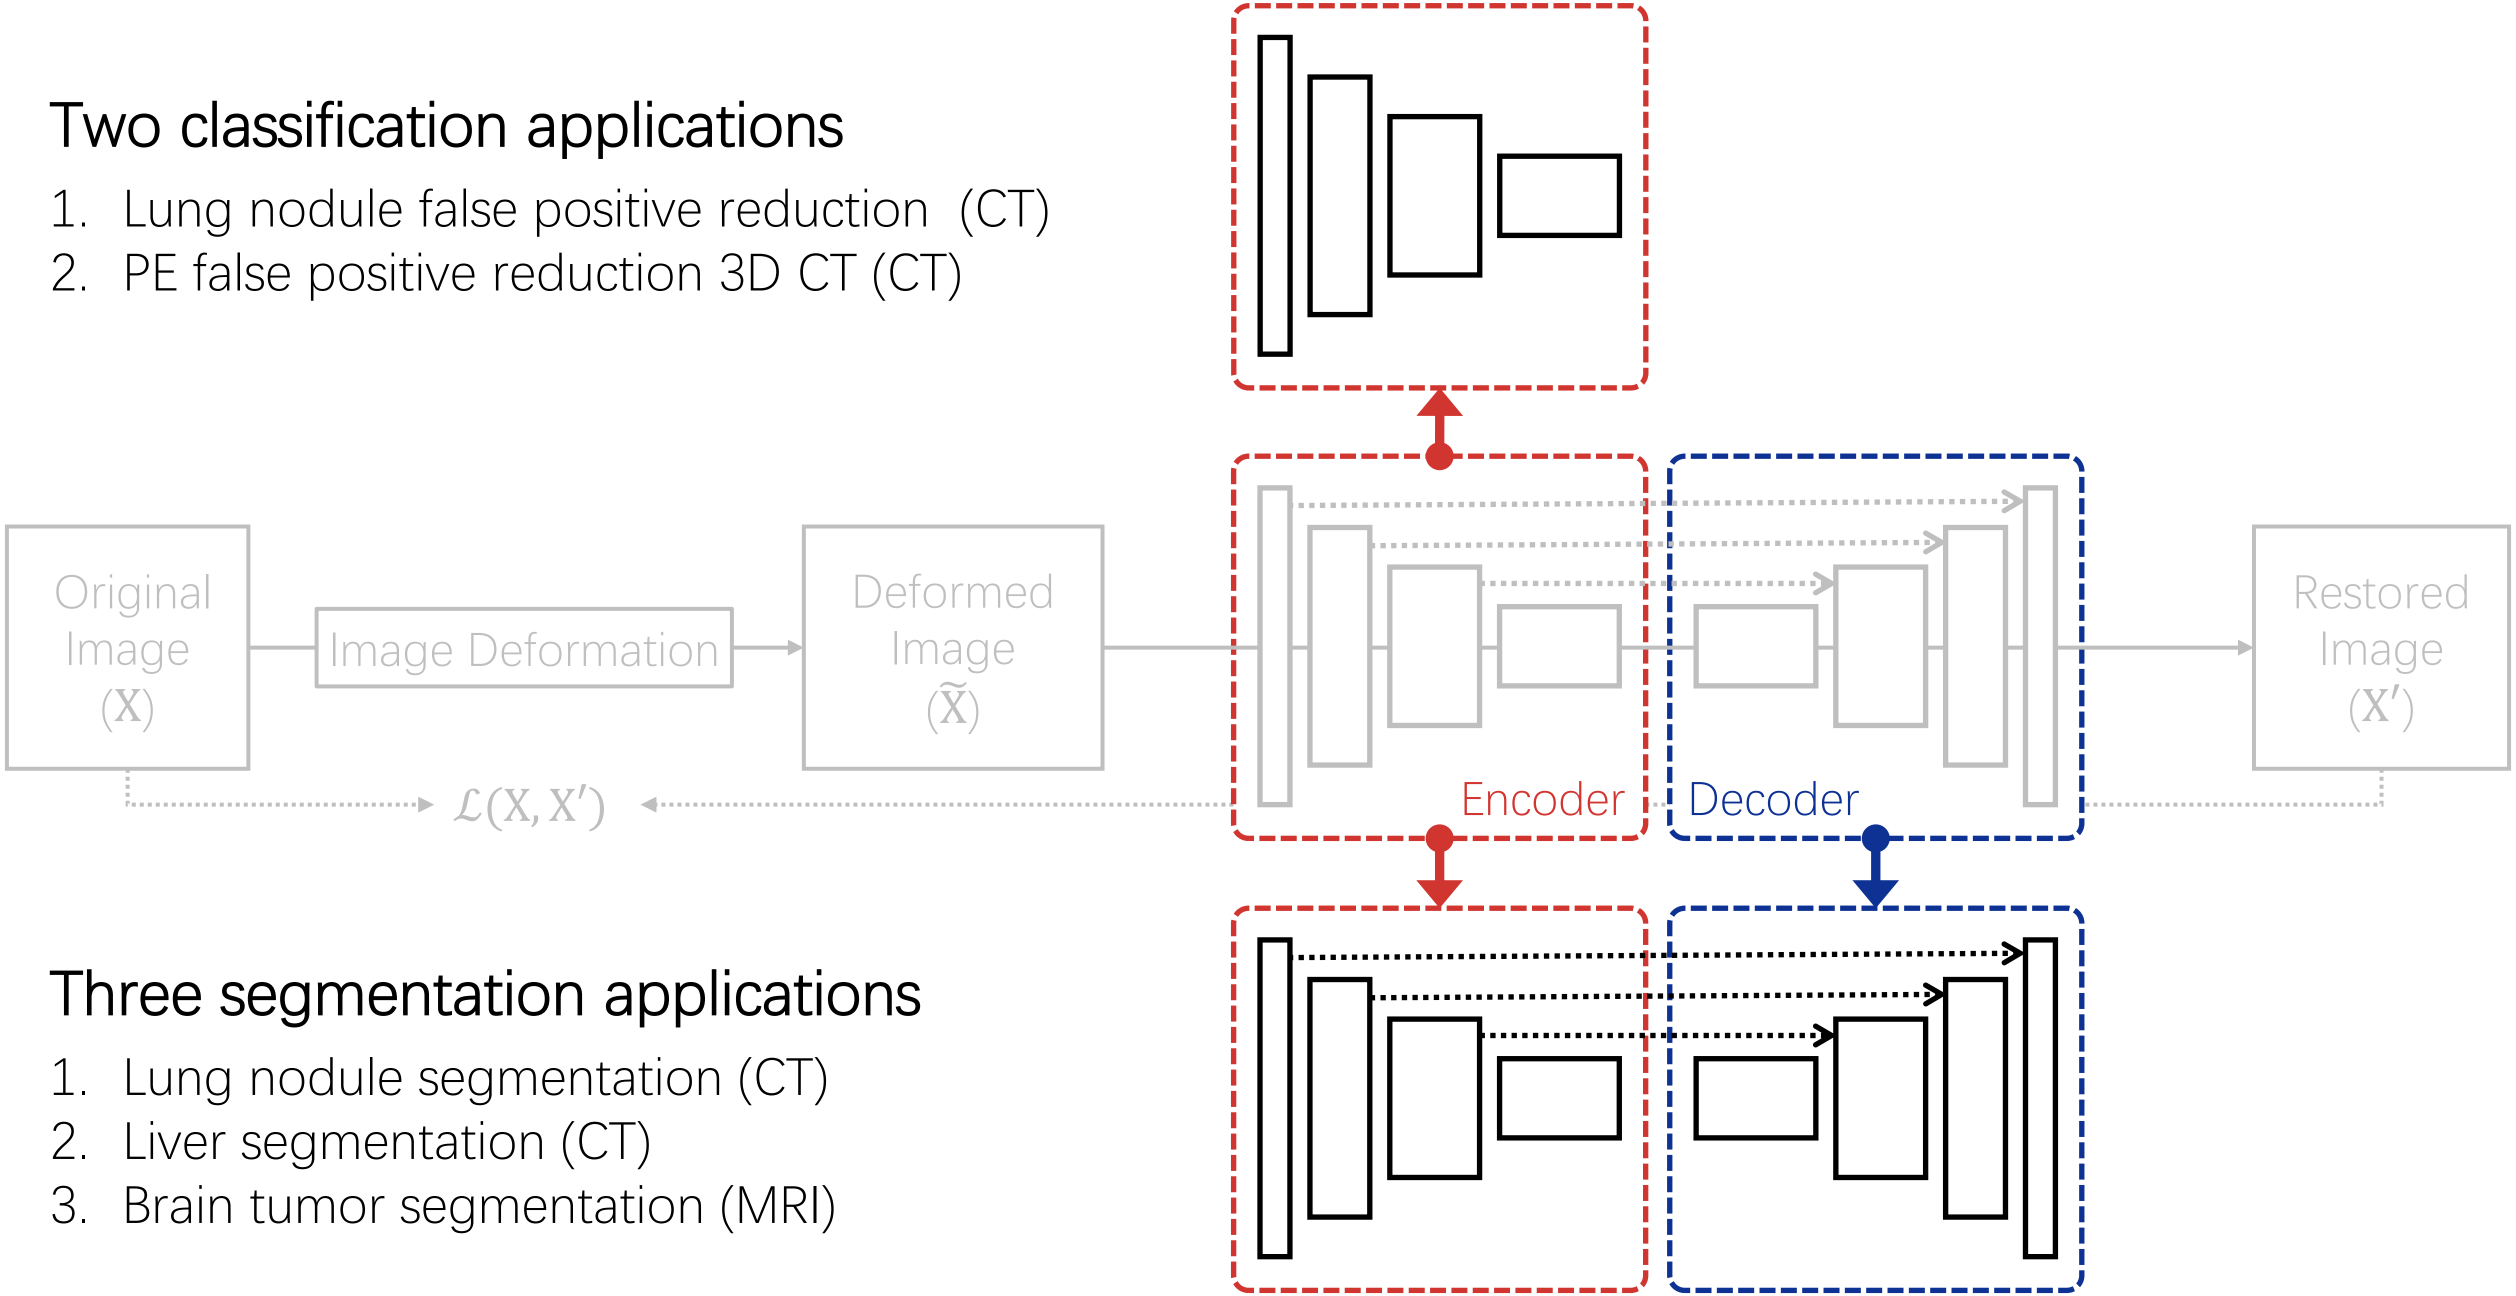

Three segmentation applications

1. Lung nodule segmentation (CT)
2. Liver segmentation (CT)
3. Brain tumor segmentation (MRI)

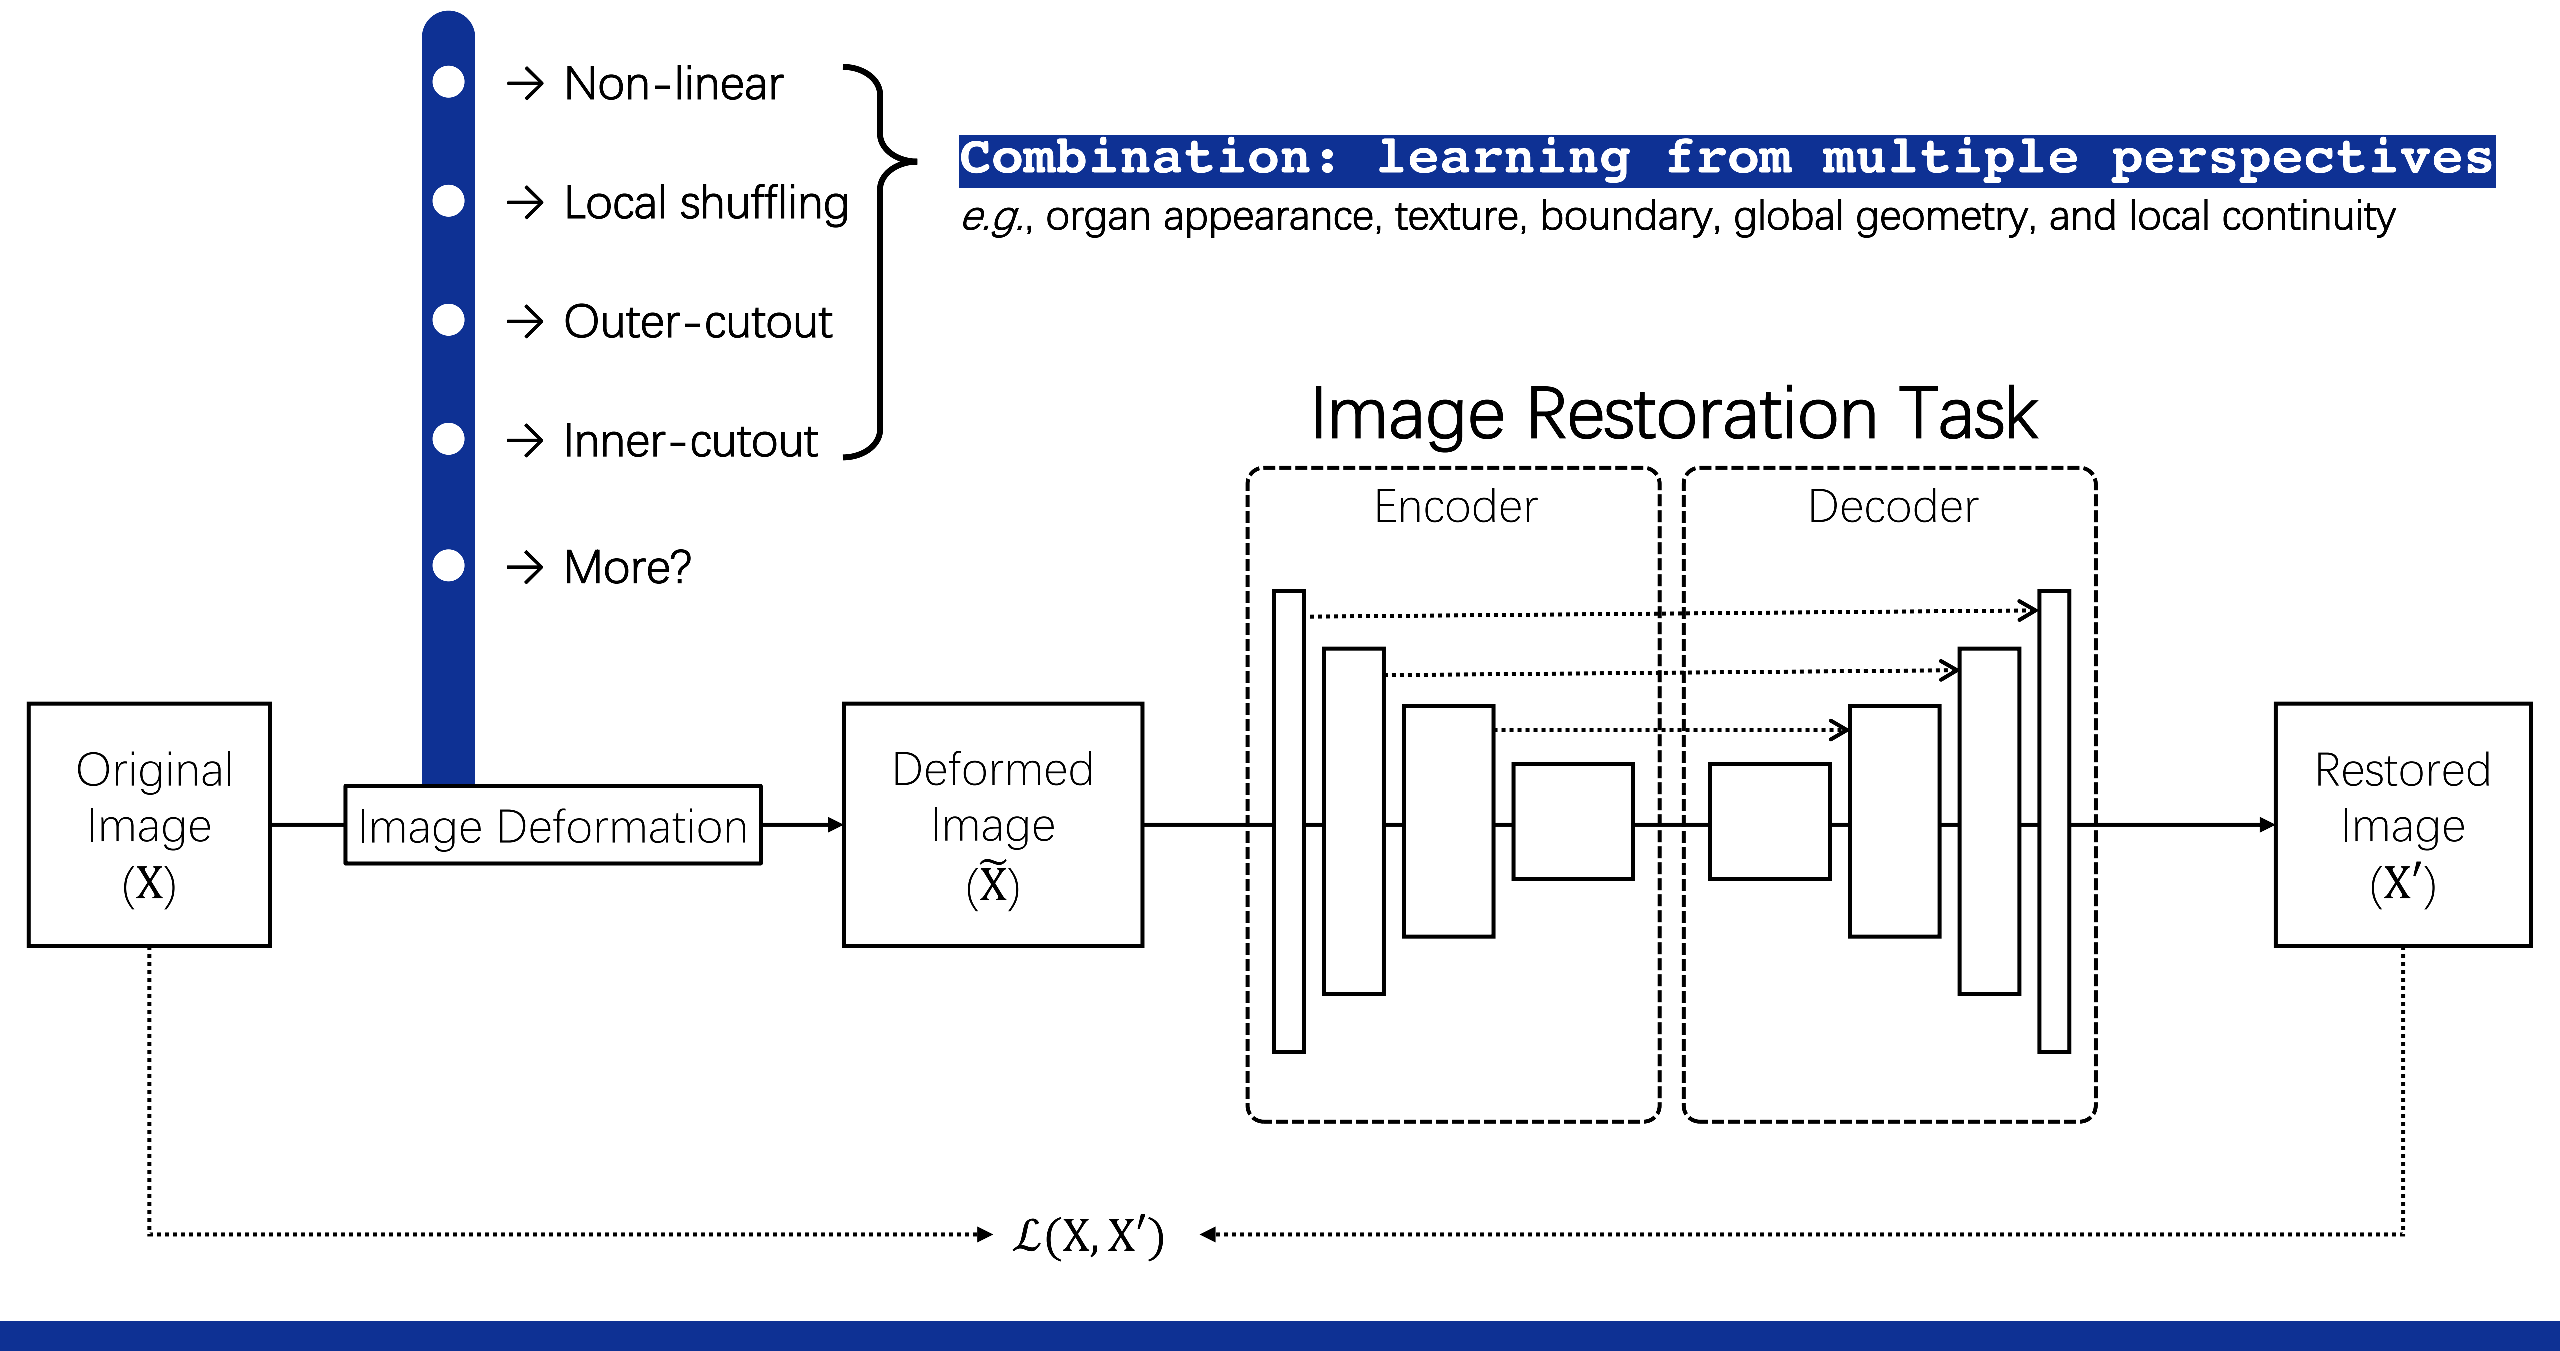

→ Non-linear
→ Local shuffling
→ Outer-cutout
→ Inner-cutout
→ More?
Combination: learning from multiple perspectives
e.g., organ appearance, texture, boundary, global geometry, and local continuity
Image Restoration Task
Encoder
Decoder
Original Image (X)
Image Deformation
Deformed Image ($\tilde{X}$)
Restored Image (X′)
$\mathcal{L}(X, X')$

Ablation study: The combined learning scheme exceeds each individual

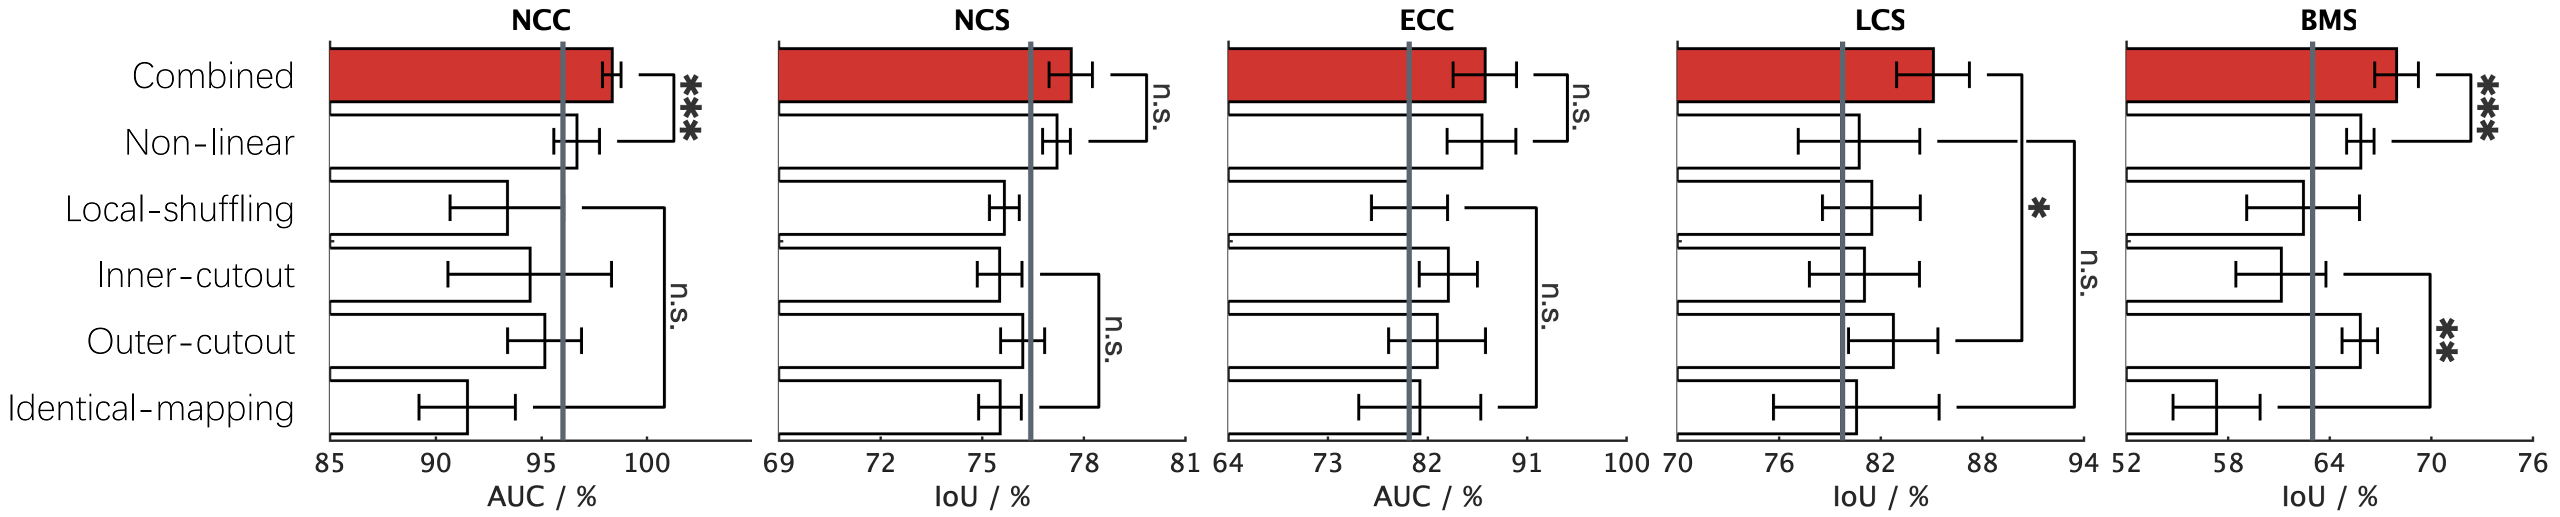

Two classification applications

1. Lung nodule false positive reduction (NCC)
2. PE false positive reduction 3D CT (CT)

Three segmentation applications

1. Lung nodule segmentation (CT)
2. Liver segmentation (CT)
3. Brain tumor segmentation (MRI)

Image restoration is promising to help

deep models learn common visual representation

Result I: Models Genesis outperform 3D models learning from scratch

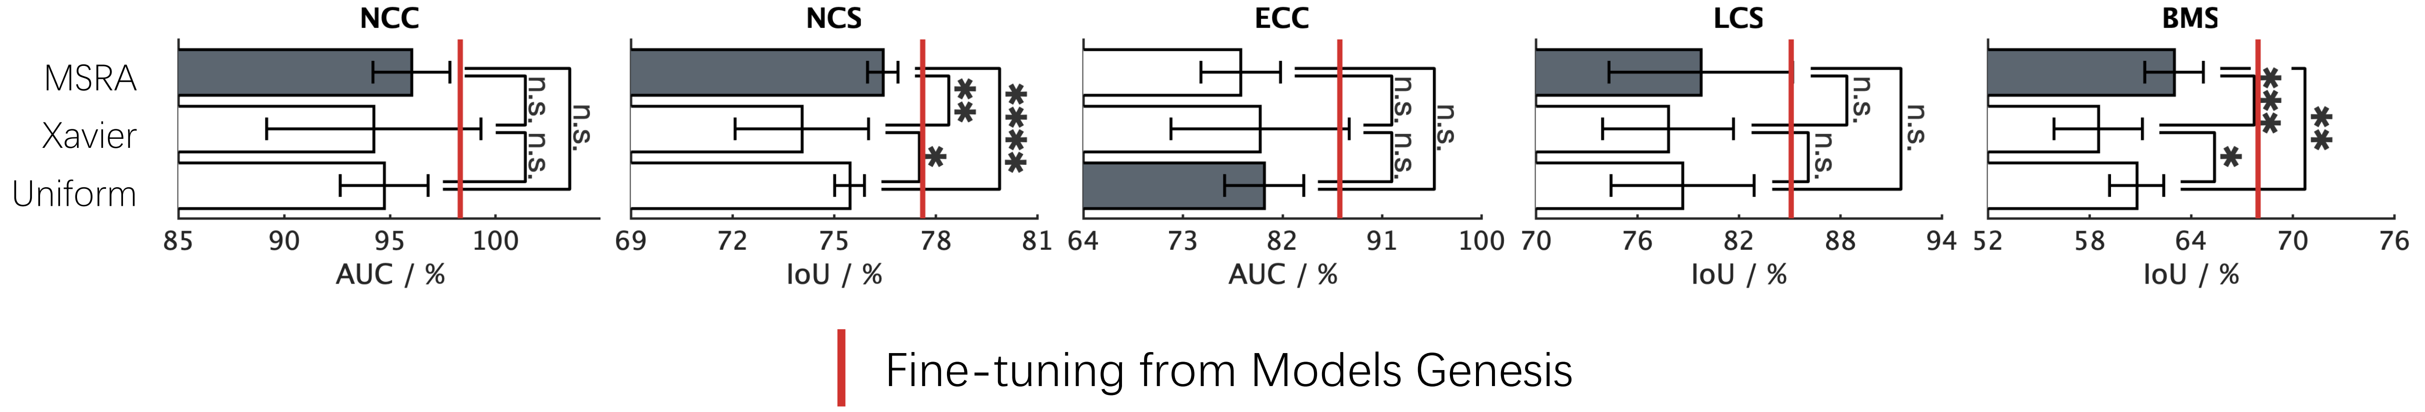

Two classification applications

1. Lung nodule false positive reduction (NCC)
2. PE false positive reduction 3D CT (CT)

Three segmentation applications

1. Lung nodule segmentation (CT)
2. Liver segmentation (CT)
3. Brain tumor segmentation (MRI)

Result II: Models Genesis surpass existing pre-trained 3D models

Proxy task	Approach \ Target task	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
Full-supervision	I3D (Carreira and Zisserman, 2017)	98.26±0.27	71.31±0.37	80.55±1.11	69.82±4.95	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-supervision	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
	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] The winner in LUNA (2016) holds an official score of 0.968 vs. 0.971 (ours)
[2] Wu et al. (2018) holds a Dice of 74.05% vs. 75.86%±0.90% (ours)
[3] Zhou et al. (2017) holds an AUC of 87.06% vs. 87.20%±2.87% (ours)
[4] The winner in LiTS (2017) with post-processing holds a Dice of 96.60% vs. 93.19%±0.46% (ours without post-processing)
[5] We have only investigated the transfer learning from CT to MR Flair image domain, so the results are not submitted to BraTS 2018.
[†] Genesis Chest CT is slightly outperformed by MedicalNet in LCS because the latter was supervised pre-trained on the LiTS dataset in the proxy task.

Result III: Models Genesis reduce annotation efforts by at least 30%

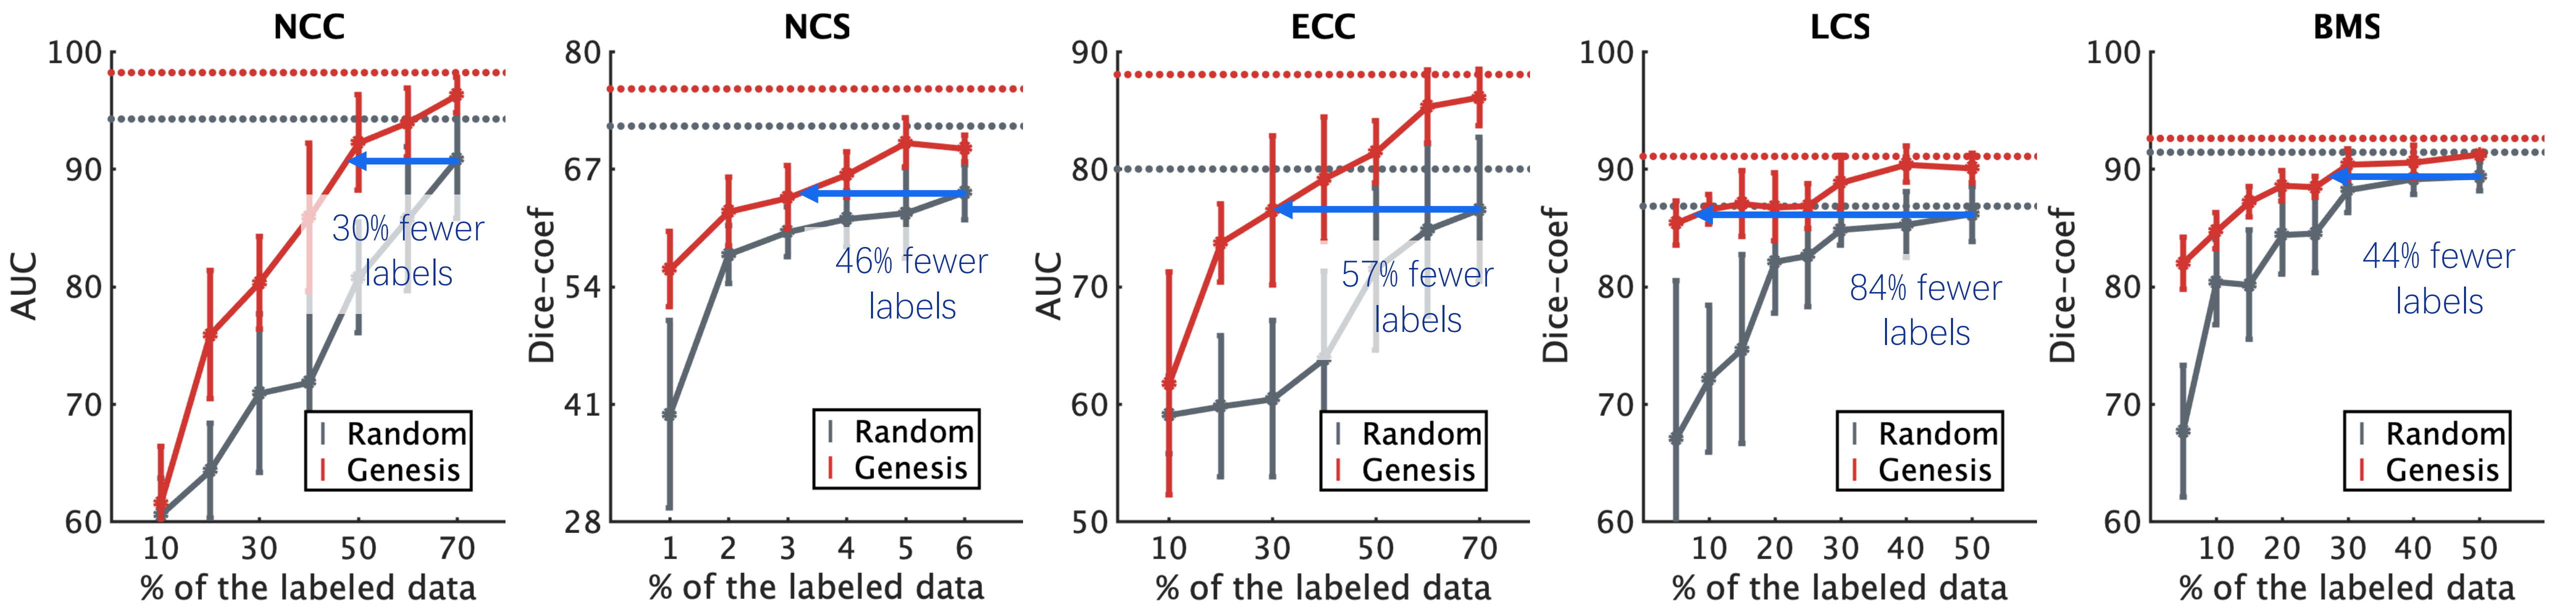

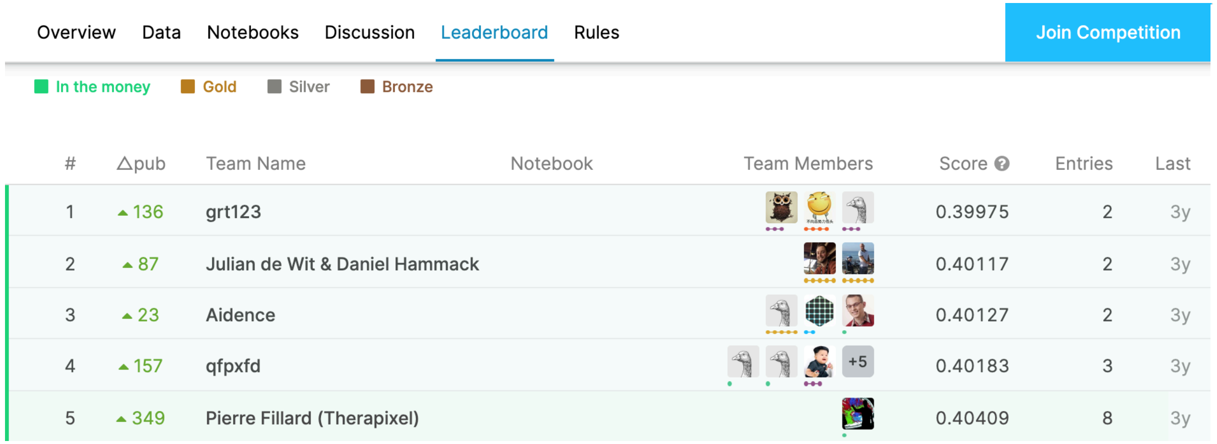

#	△pub	Team Name	Notebook	Team Members	Score	Entries	Last
1	▲ 136	grt123			0.39975	2	3y
2	▲ 87	Julian de Wit & Daniel Hammack			0.40117	2	3y
3	▲ 23	Aidence			0.40127	2	3y
4	▲ 157	qfpxfd		+5	0.40183	3	3y
5	▲ 349	Pierre Fillard (Therapixel)			0.40409	8	3y

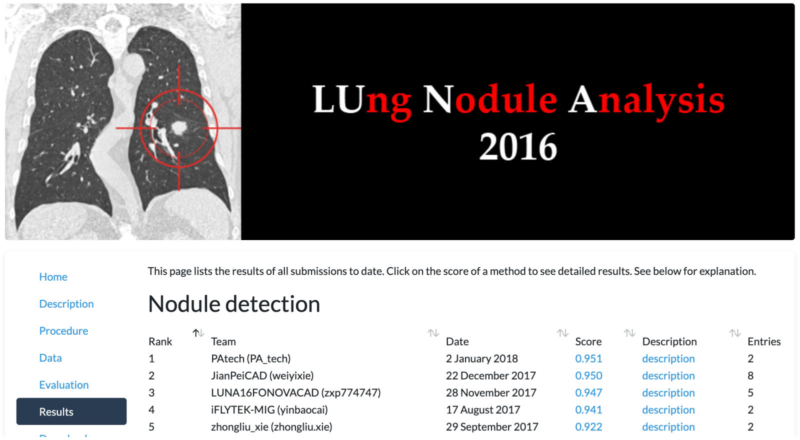

Rank	Team	Date	Score	Description	Entries
1	PAtech (PA_tech)	2 January 2018	0.951	description	2
2	JianPeiCAD (weiyixie)	22 December 2017	0.950	description	8
3	LUNA16FONOVACAD (zxp774747)	28 November 2017	0.947	description	5
4	iFLYTEK-MIG (yinbaocai)	17 August 2017	0.941	description	2
5	zhongliu_xie (zhongliu.xie)	29 September 2017	0.922	description	2

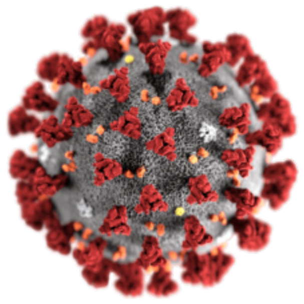

COVID-19?

Call for Papers

IEEE Transactions on Medical Imaging Special Issue on

Annotation-Efficient Deep Learning for Medical Imaging

https://ieee-tmi.org/Special_Issue_CFP_DL4MI.pdf

Important Dates

Paper submission deadline:	July 1, 2020
First reviews due:	September 1, 2020
Revised manuscript due:	November 1, 2020
Final decision:	December 1, 2020
Camera ready version:	December 15, 2020

Guest Editors

Nima Tajbakhsh, PhD	VoxelCloud, Inc.	ntajbakhsh@voxelcloud.io
Jianming Liang, PhD	Arizona State University	jianming.liang@asu.edu
Holger Roth, PhD	Nvidia, Inc.	hroth@nvidia.com
Demetri Terzopoulos, PhD	University of California, Los Angeles	dt@cs.ucla.edu

Result IV: Models Genesis consistently top any 2D/2.5D approaches

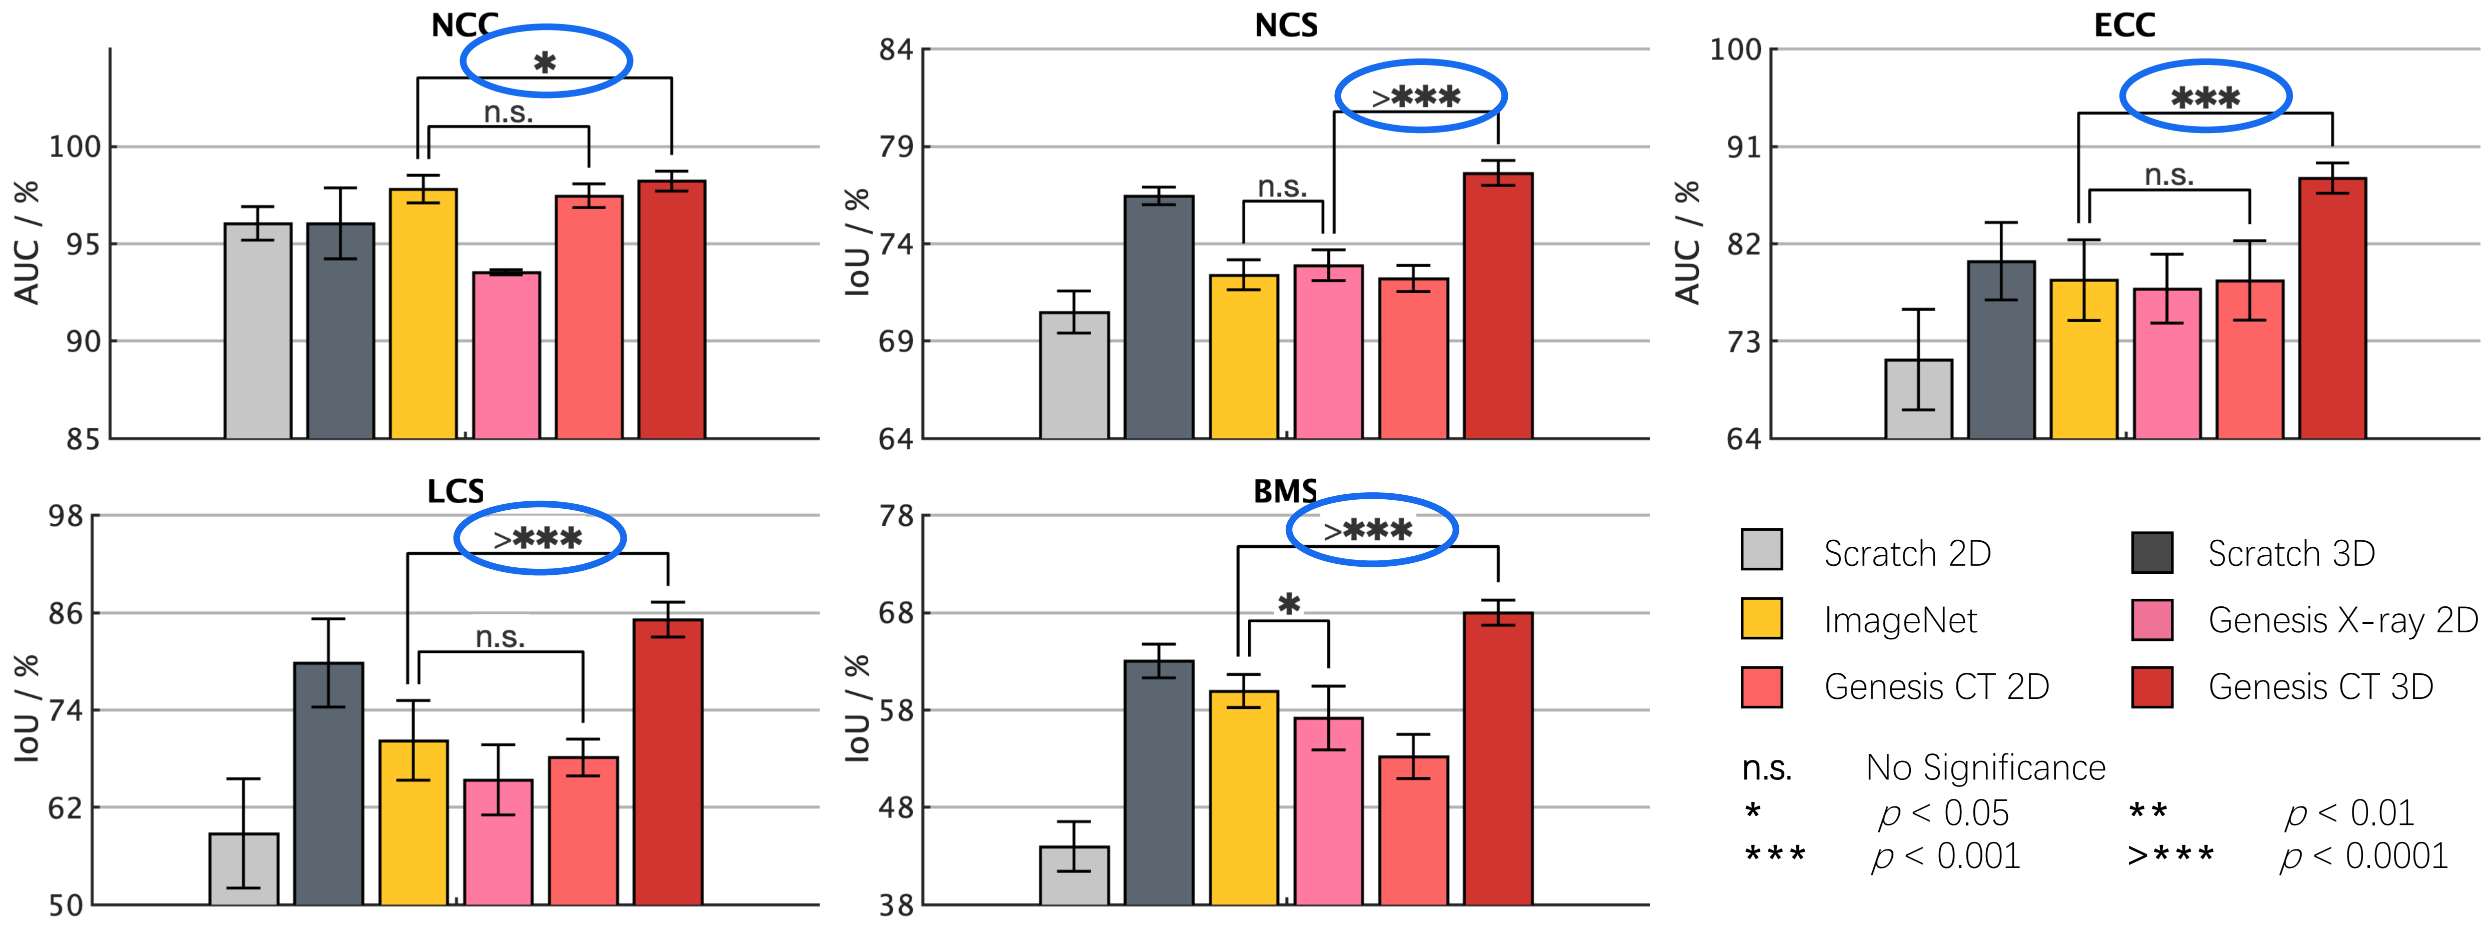

Result IV: Models Genesis consistently top any 2D/2.5D approaches

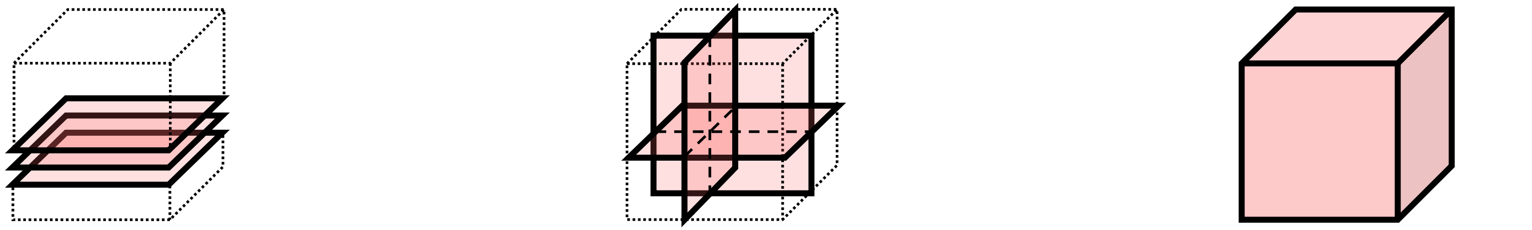

2D slice-based | 2.5D orthogonal | 3D volume-based

Task: `NCC`	Random	ImageNet	Genesis
2D slice-based input	96.03±0.86	**97.79±0.71**	97.45±0.61
2.5D orthogonal input	95.76±1.05	**97.24±1.01**	97.07±0.92
3D volume-based input	96.03±1.82	n/a	**98.34±0.44**
Task: `ECC`	Random	ImageNet	Genesis
2D slice-based input	60.33±8.61	62.57±8.04	**62.84±8.78**
2.5D orthogonal input	71.27±4.64	**78.61±3.73**	78.58±3.67
3D volume-based input	80.36±3.58	n/a	**88.04±1.40**

3D problems should be solved in 3D, directly but properly.

Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis

We offer a set of powerful pre-trained 3D models, concluding that

1. Models Genesis outperform 3D models trained from scratch
2. Models Genesis surpass existing pre-trained 3D models
3. Models Genesis reduce annotation efforts by at least 30%
4. Models Genesis consistently top any 2D/2.5D approaches

Genesis Chest CT

Genesis X-ray

We hope that our collective efforts will lead to the Holy Grail of Models Genesis, effective across diseases, organs, and modalities.

“通用表征学习的好处在于，单个任务的数据量不大，难以训练一个好的模型。如果将所有任务放在一起，就会有更多的数据，进而更好地提升模型的性能。“我们希望可以学到一个通用性的表达，对所有的任务都能适用。””

——周少华博士

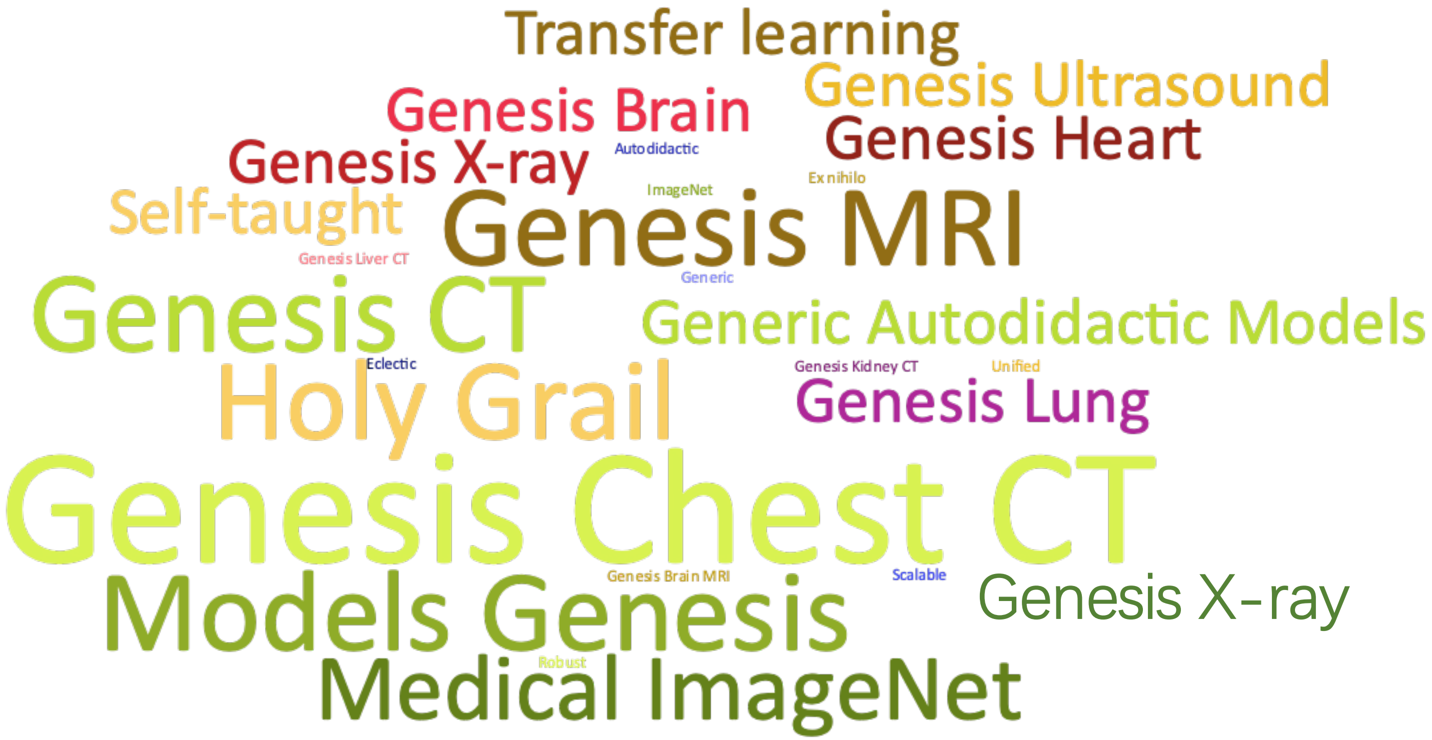

MrGiovanni / ModelsGenesis

Unwatch 13 | Star 211 | Fork 56

Code | Issues 0 | Pull requests 0 | Actions | Projects 0 | Wiki | Security | Insights | Settings

Official Keras&PyTorch Implementation and Pre-trained Models for Models Genesis - MICCAI 2019

Edit

Paper

This repository provides the official implementation of training Models Genesis as well as the usage of the pre-trained Models Genesis in the following paper:

Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis
Zongwei Zhou[1], Vatsal Sodha[1], Md Mahfuzur Rahman Siddiquee[1],
Ruibin Feng[1], Nima Tajbakhsh[1], Michael B. Gotway[2], and Jianming Liang[1]
[1] Arizona State University, [2] Mayo Clinic
International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI), 2019
Young Scientist Award
paper | code | slides | poster | talk (YouTube, YouKu) | blog

Fine-tune from our pre-trained Models Genesis

Models Genesis: PyTorch version available Inbox ×

Zongwei Zhou <zzhou82@asu.edu>
to Zongwei, bcc:

We make the development of Models Genesis open science and invite researchers around the world to contribute to this effort.

Greetings,

Today we release **Models Genesis** official implementation in PyTorch.
The pre-trained weights, for both *Keras* and *PyTorch*, are now publicly available as well.

Download from Google Drive: https://drive.google.com/drive/folders/1H_e0PKPJSRwnAulE3XAQqfcgtK1gK8Cq?usp=sharing

Or download from Baidu Wangpan: https://pan.baidu.com/s/1Qnb1M9i0eeMZ_C416wm1bw (Pass: 72c7)

For more information, please refer to our project page.
Thank you for your interest!

Zongwei Zhou
Ph.D. Candidate | Research Assistant
College of Health Solutions, Arizona State University
13212 E Shea Blvd, Scottsdale, AZ 85259
P: 1-(480)738-2575 | E: zongweiz@asu.edu

We provide pre-trained 3D models!

Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis

Questions?

@MrGiovanni

Zongwei Zhou[1], Vatsal Sodha[1], Md Mahfuzur Rahman Siddiquee[1], Ruibin Feng[1], Nima Tajbakhsh[1], Michael B. Gotway[2], and Jianming Liang[1]

[1] Arizona State University [2] Mayo Clinic