We provide pre-trained 3D models!

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

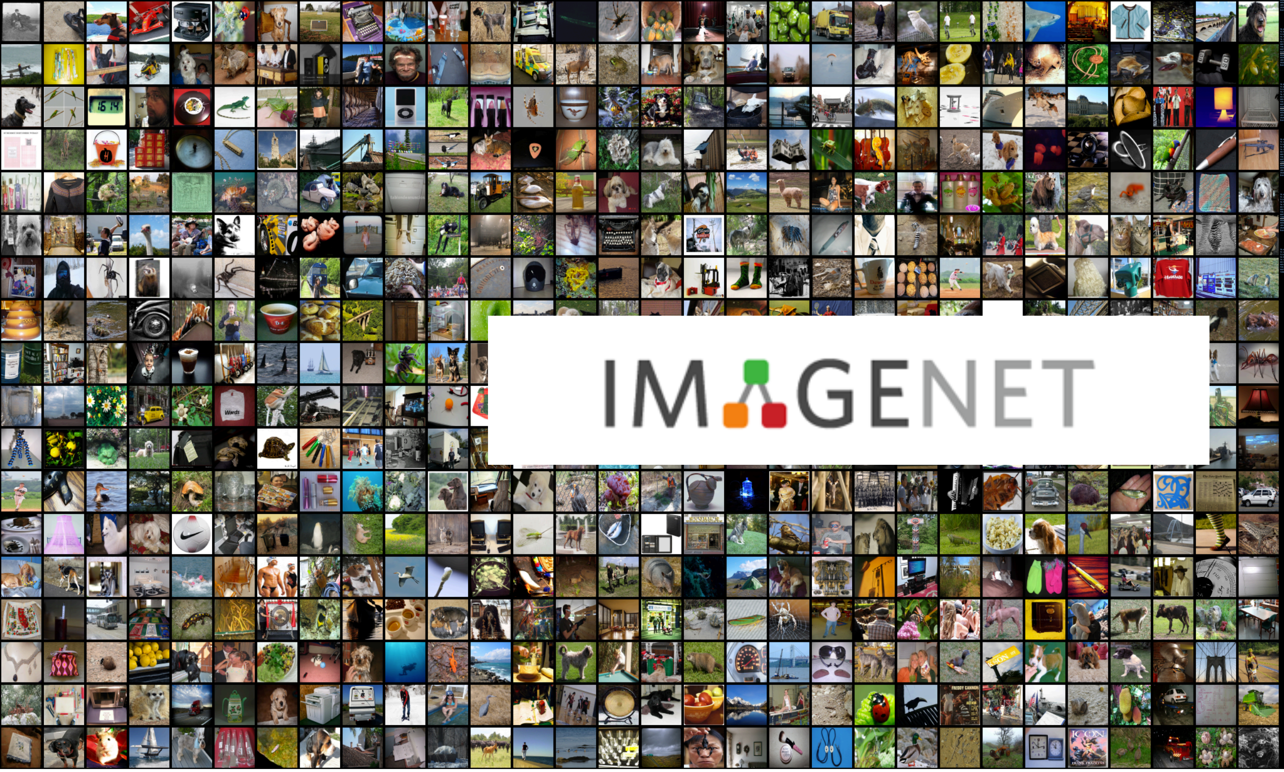
IMAGENET

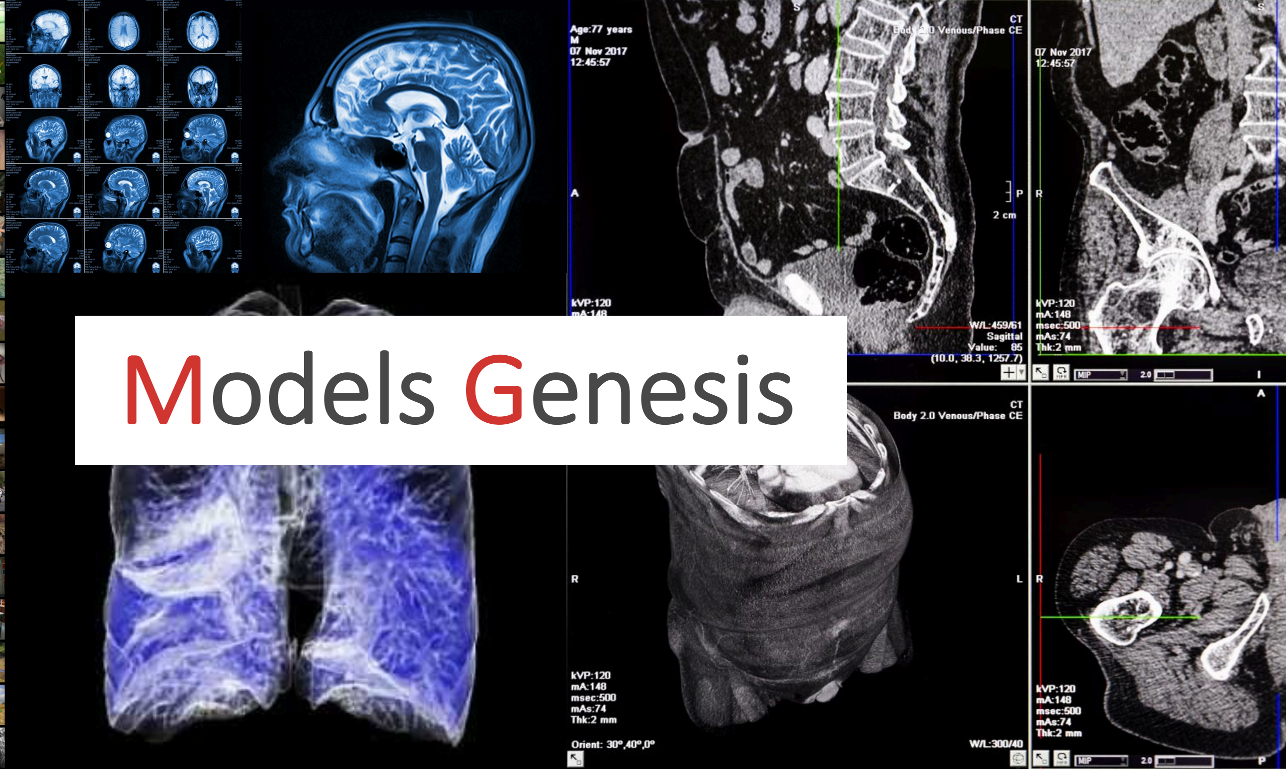
Models Genesis

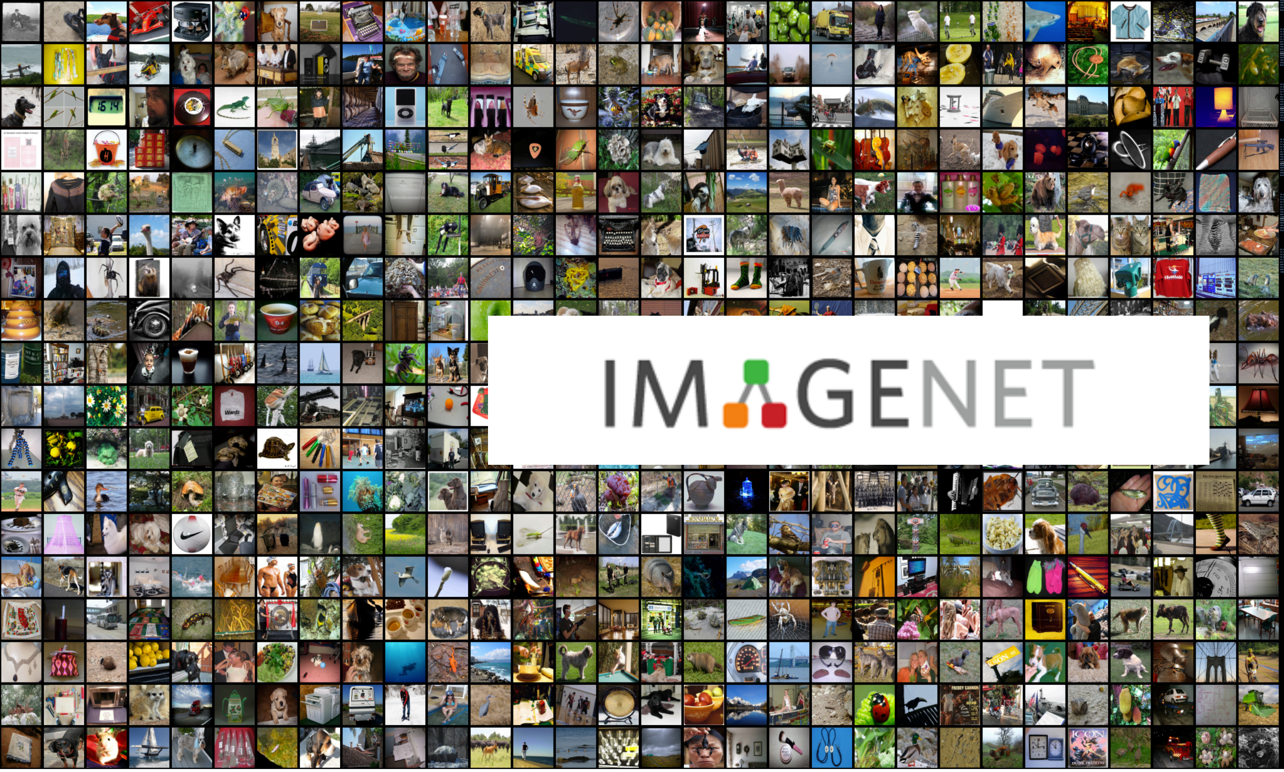

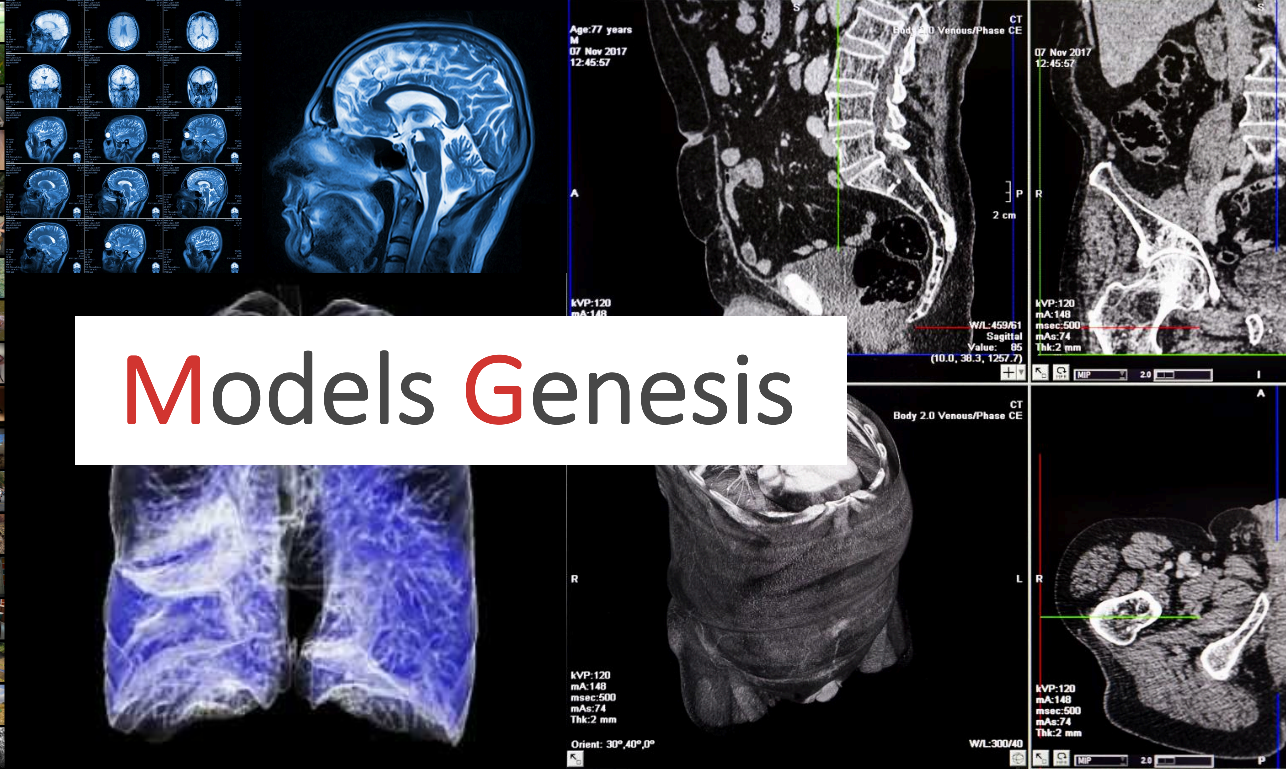

Natural images

Medical images

Transfer learning: medical images → medical images > natural images → medical images

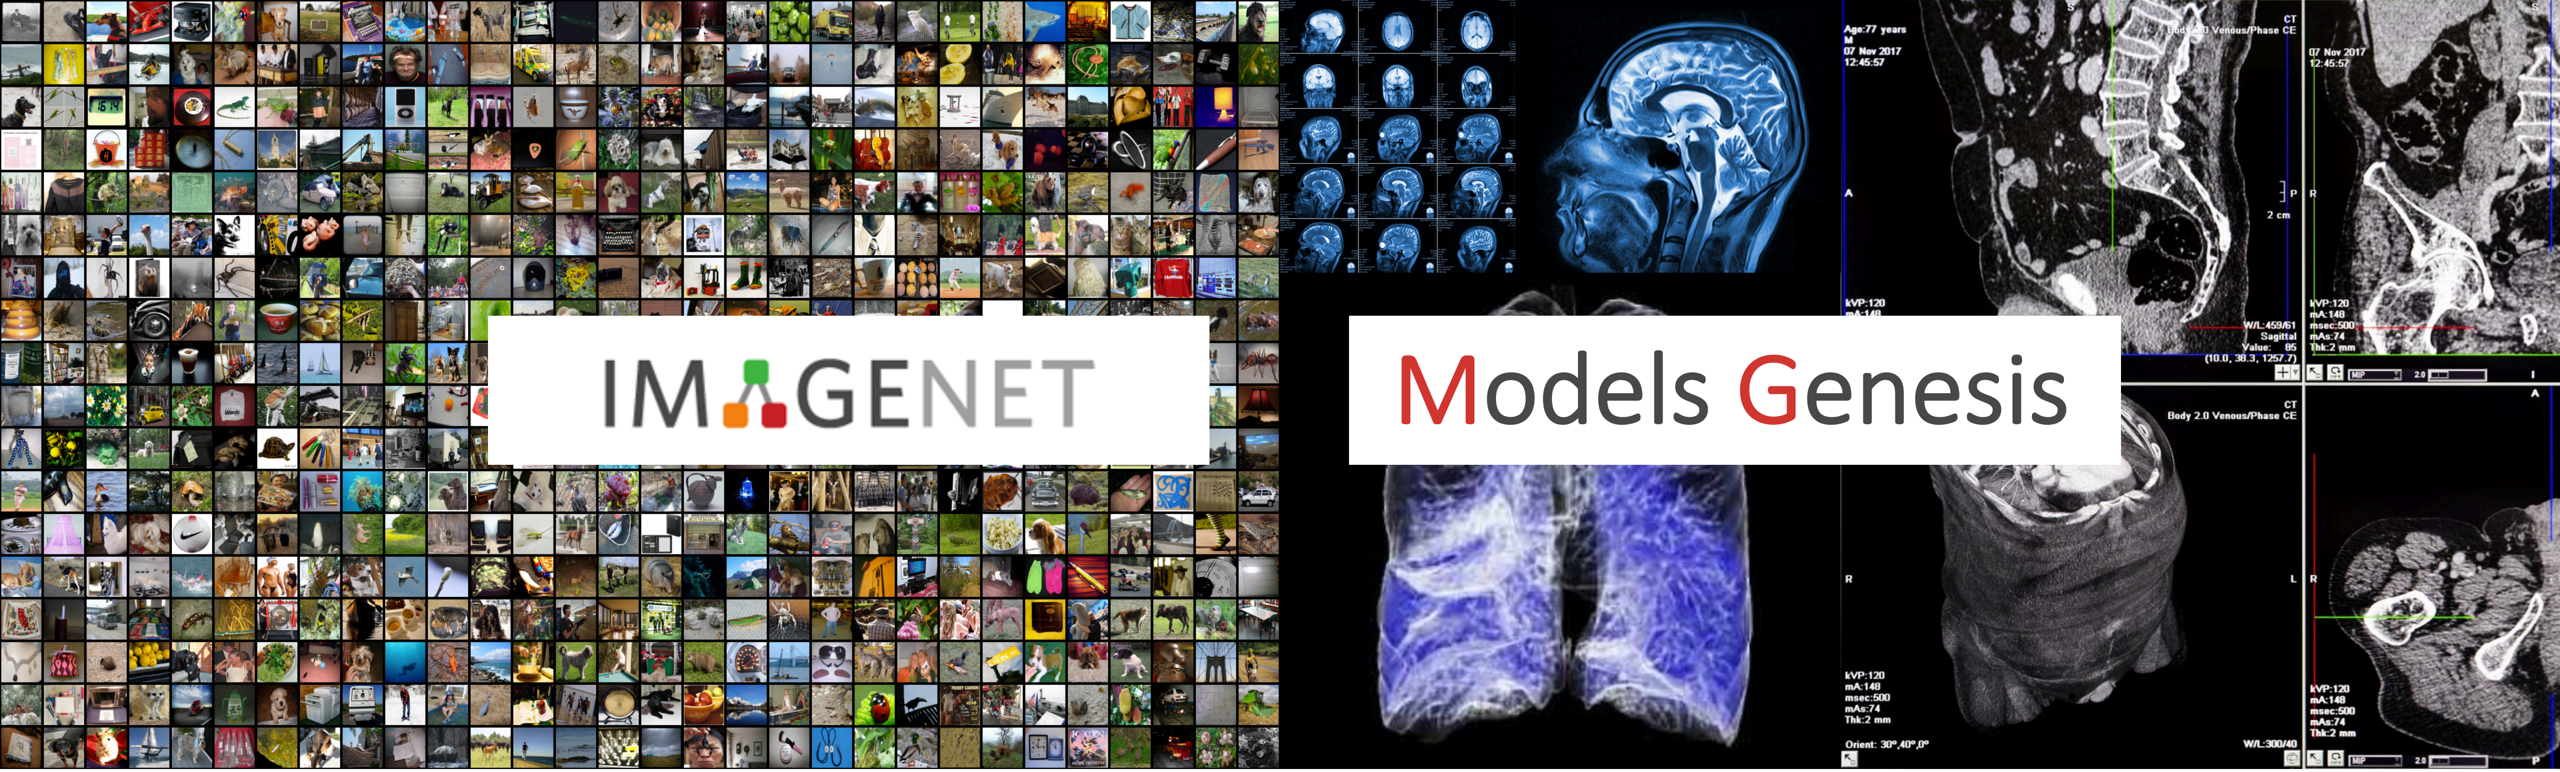

Natural images

Medical images

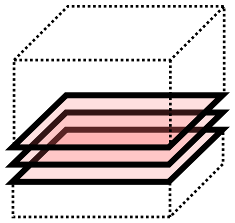

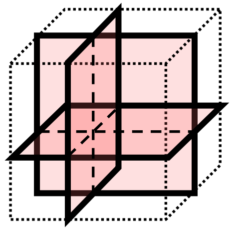

Formed in 2D

Formed in 3D

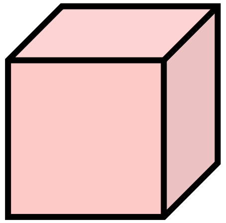

3D imaging tasks should be solved in 3D

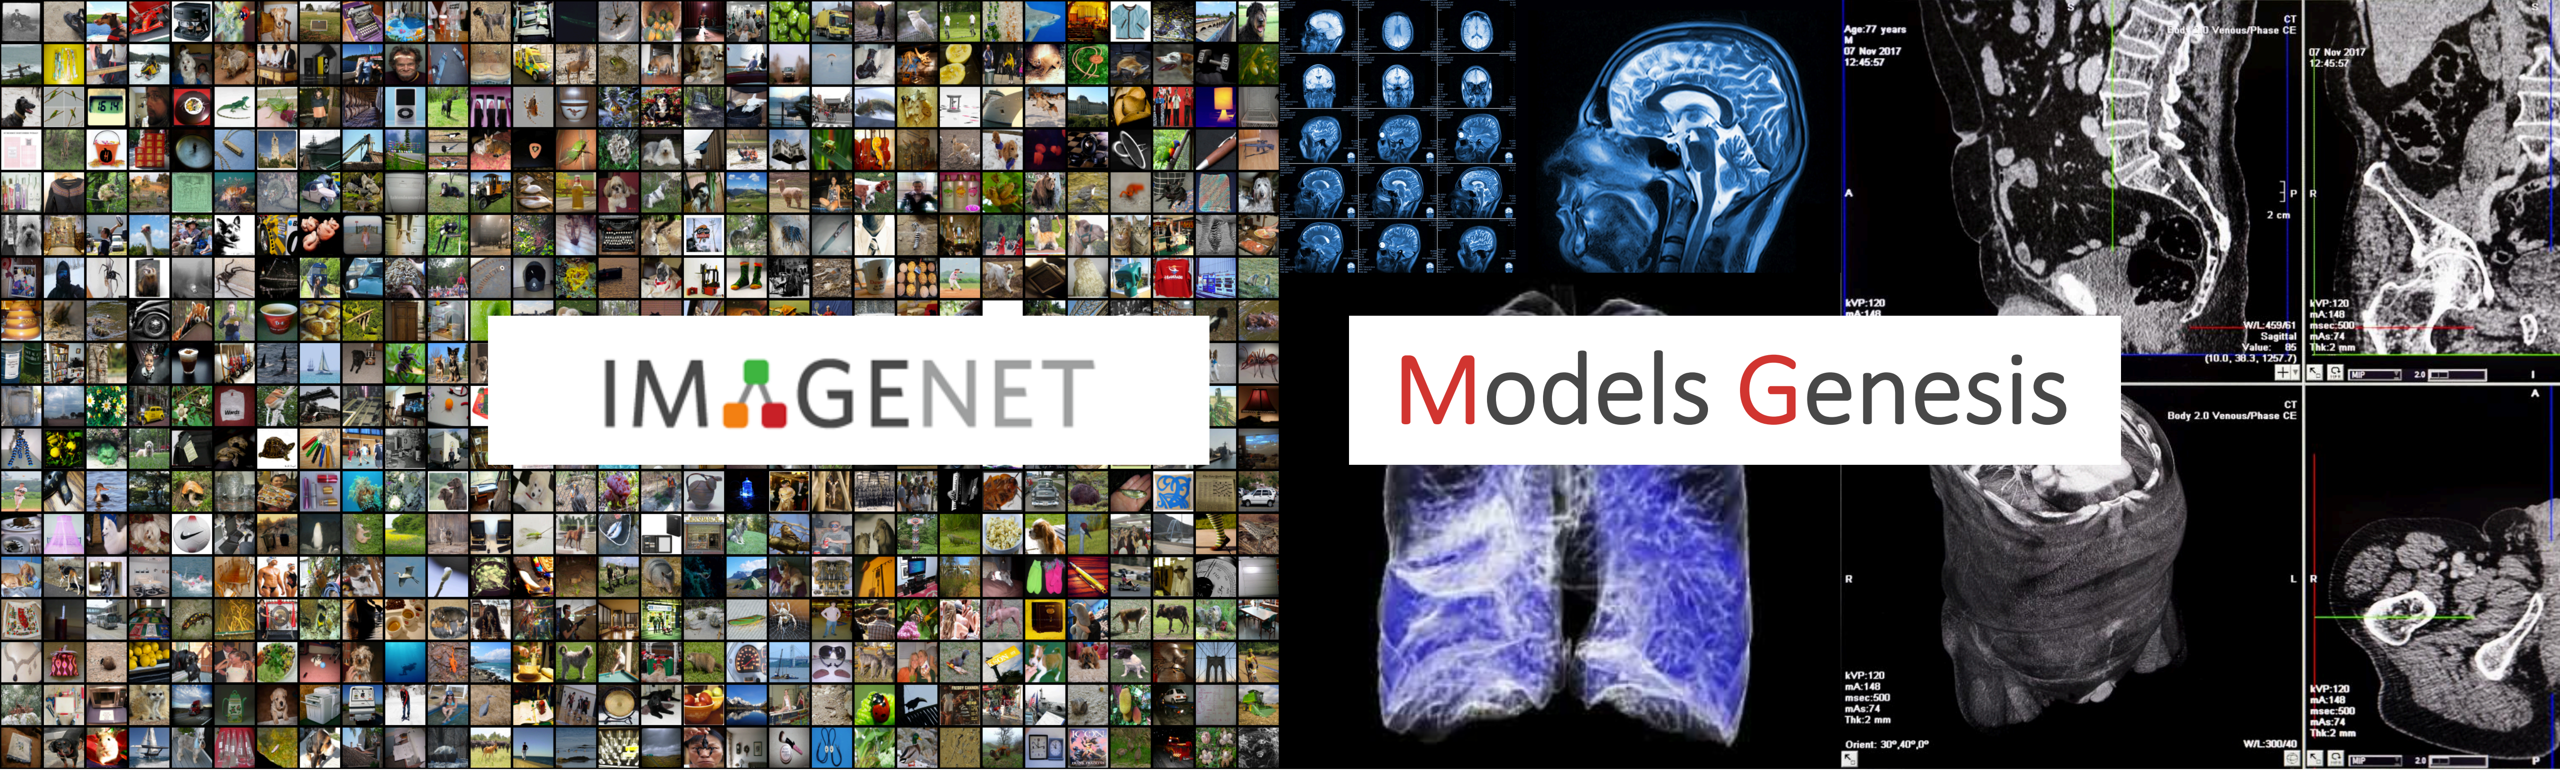

Natural images	Medical images
Formed in 2D	Formed in 3D
>14,000,000 annotation	Zero annotation

Annotating biomedical images is time consuming and demanding of costly, specialty-oriented knowledge

Image Restoration Task

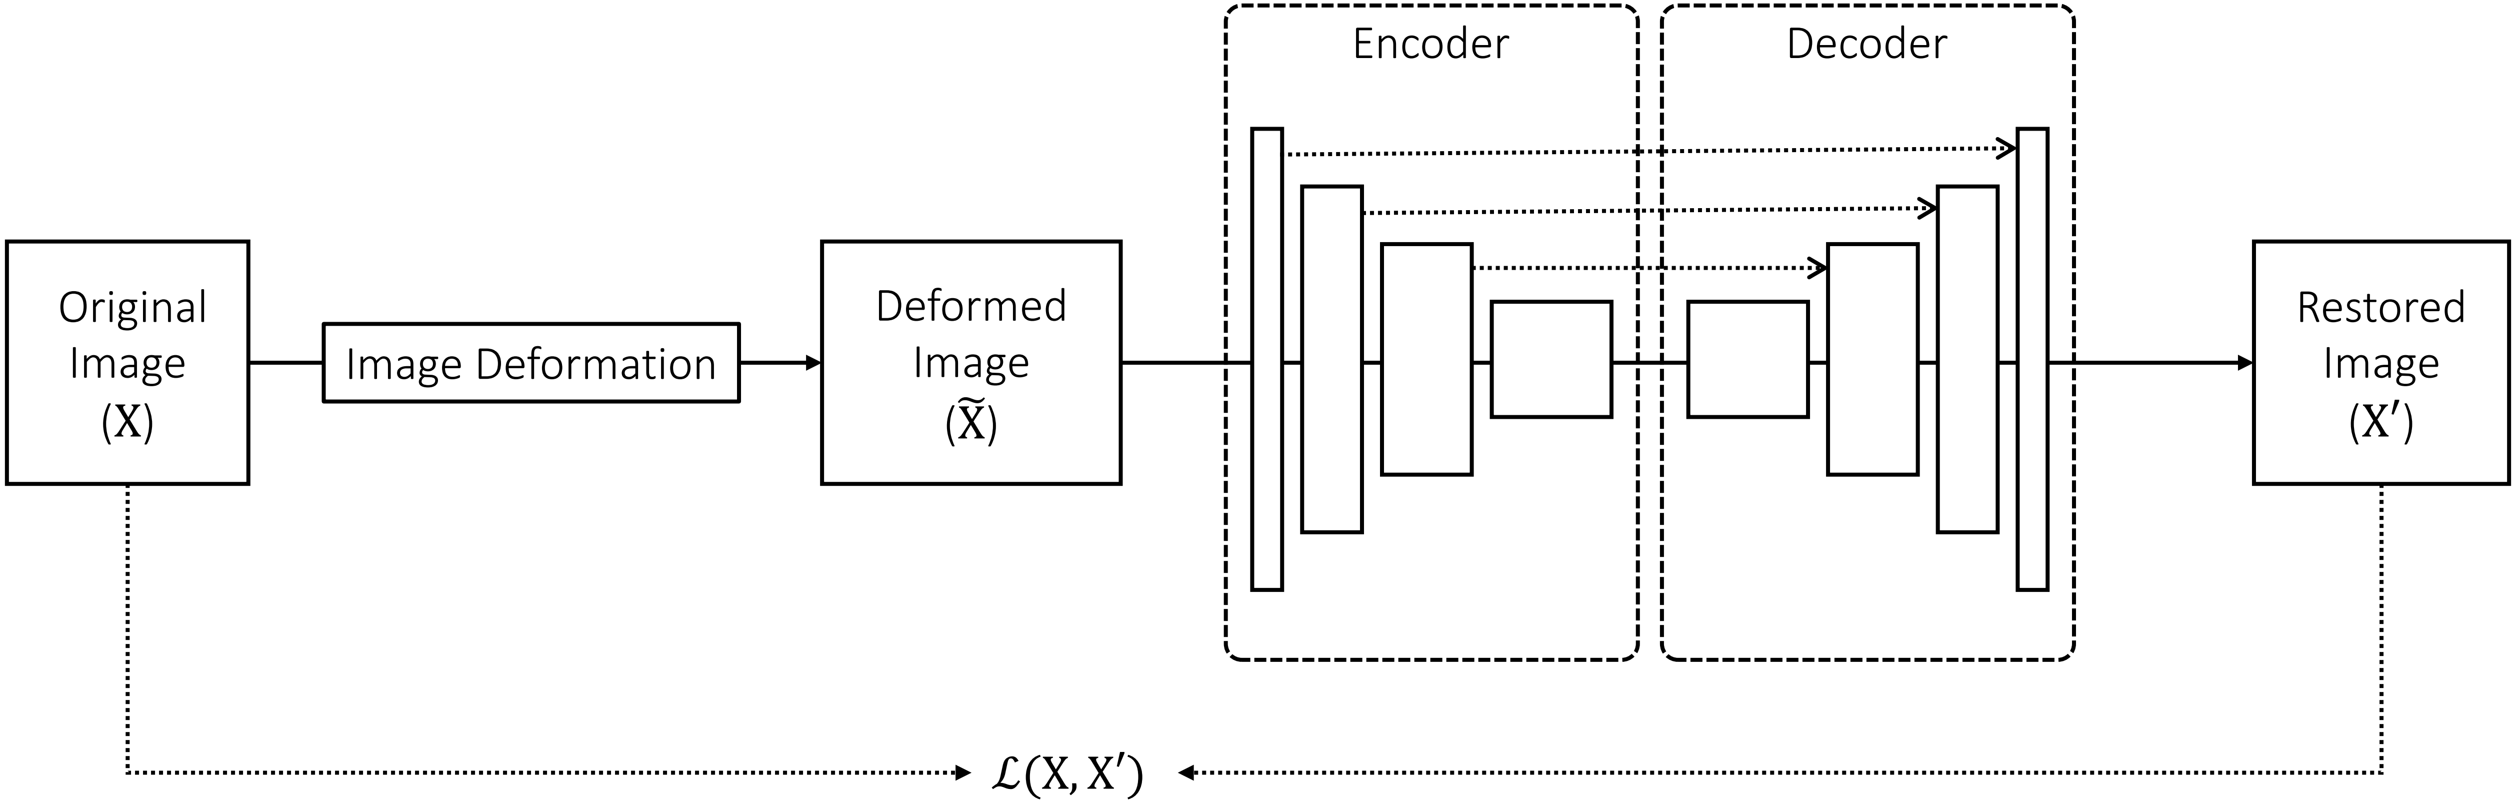

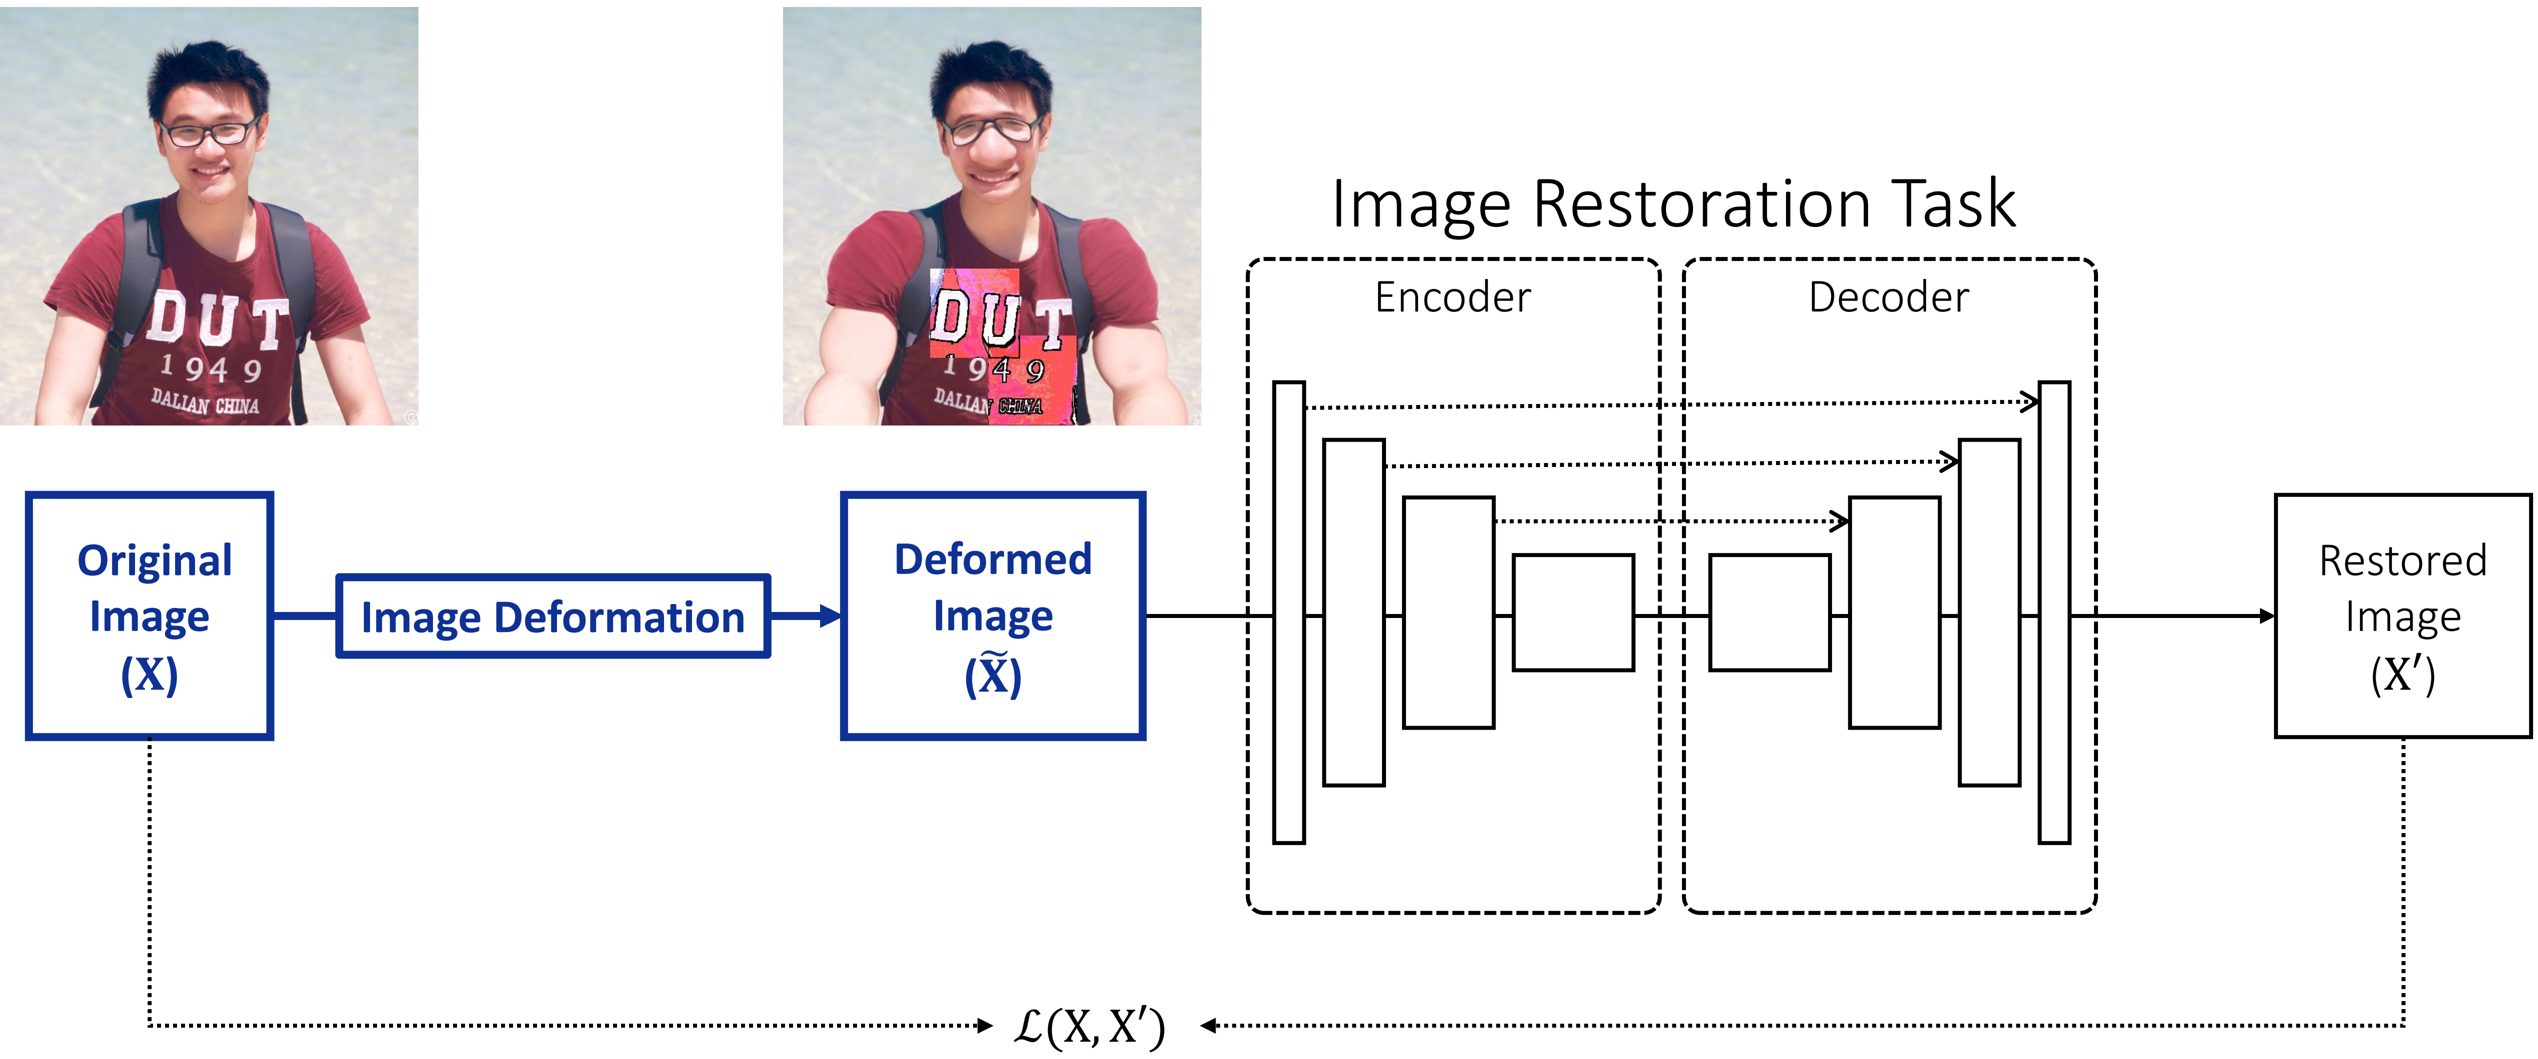
Image Restoration Task
Encoder
Decoder
Original Image (X)
Image Deformation
Deformed Image ($\tilde{X}$)
Restored Image (X′)
$\mathcal{L}(X, X')$

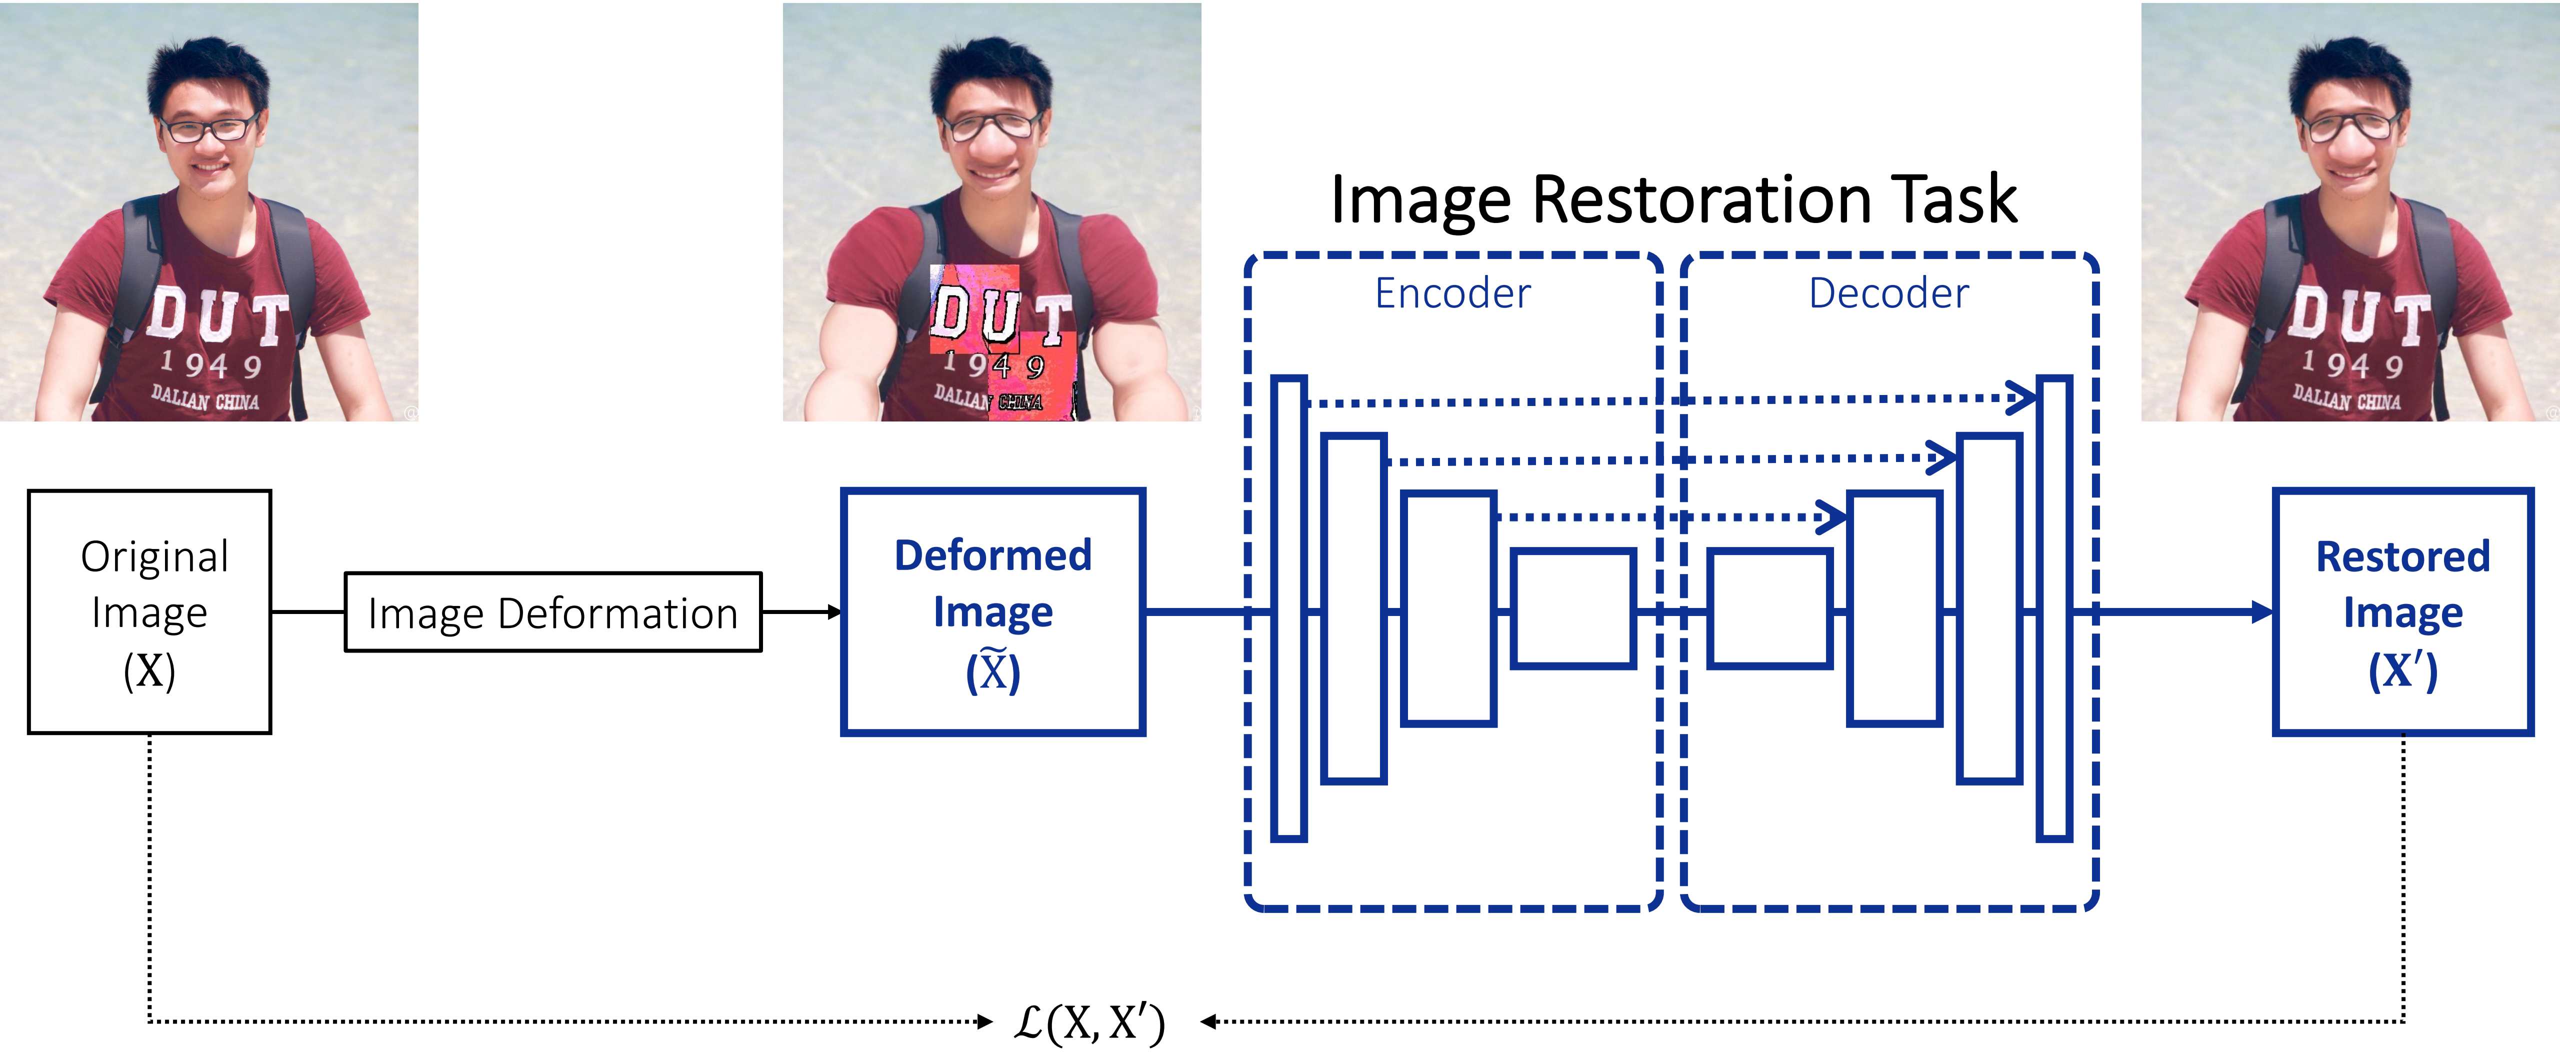
Image Restoration Task
Original Image (X)
Image Deformation
Deformed Image ($\tilde{X}$)
Encoder
Decoder
Restored Image (X′)
$\mathcal{L}(X, X')$

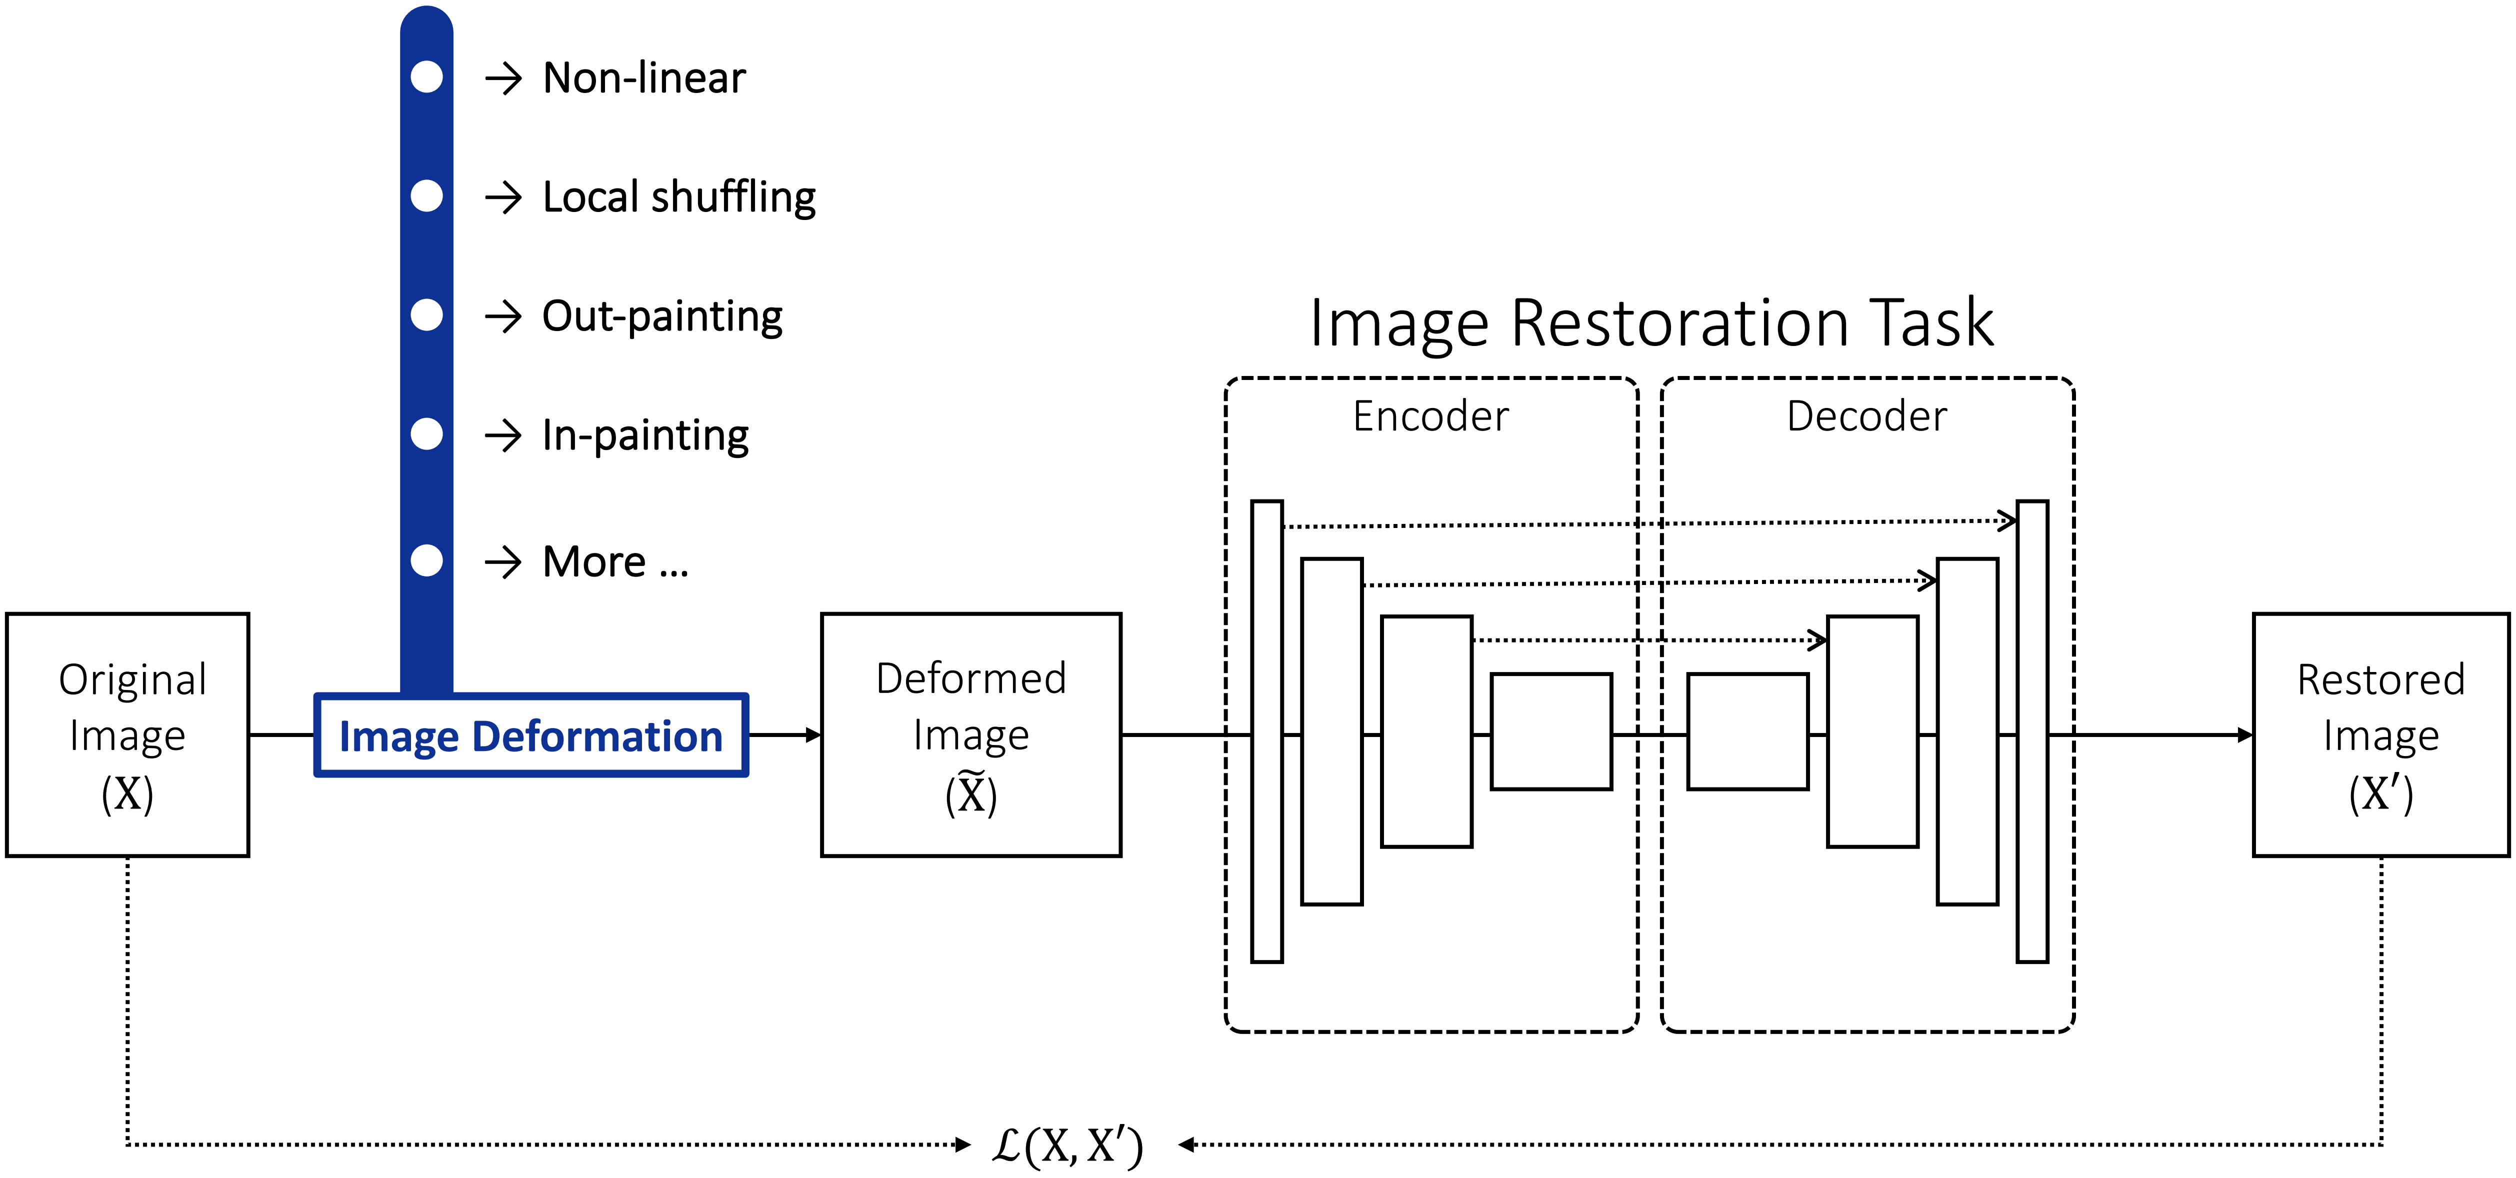
→ Non-linear
→ Local shuffling
→ Out-painting
→ In-painting
→ More ...
Original Image (X)
Image Deformation
Deformed Image ($\widetilde{X}$)
Image Restoration Task
Encoder
Decoder
Restored Image (X′)
$\mathcal{L}(X, X')$

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

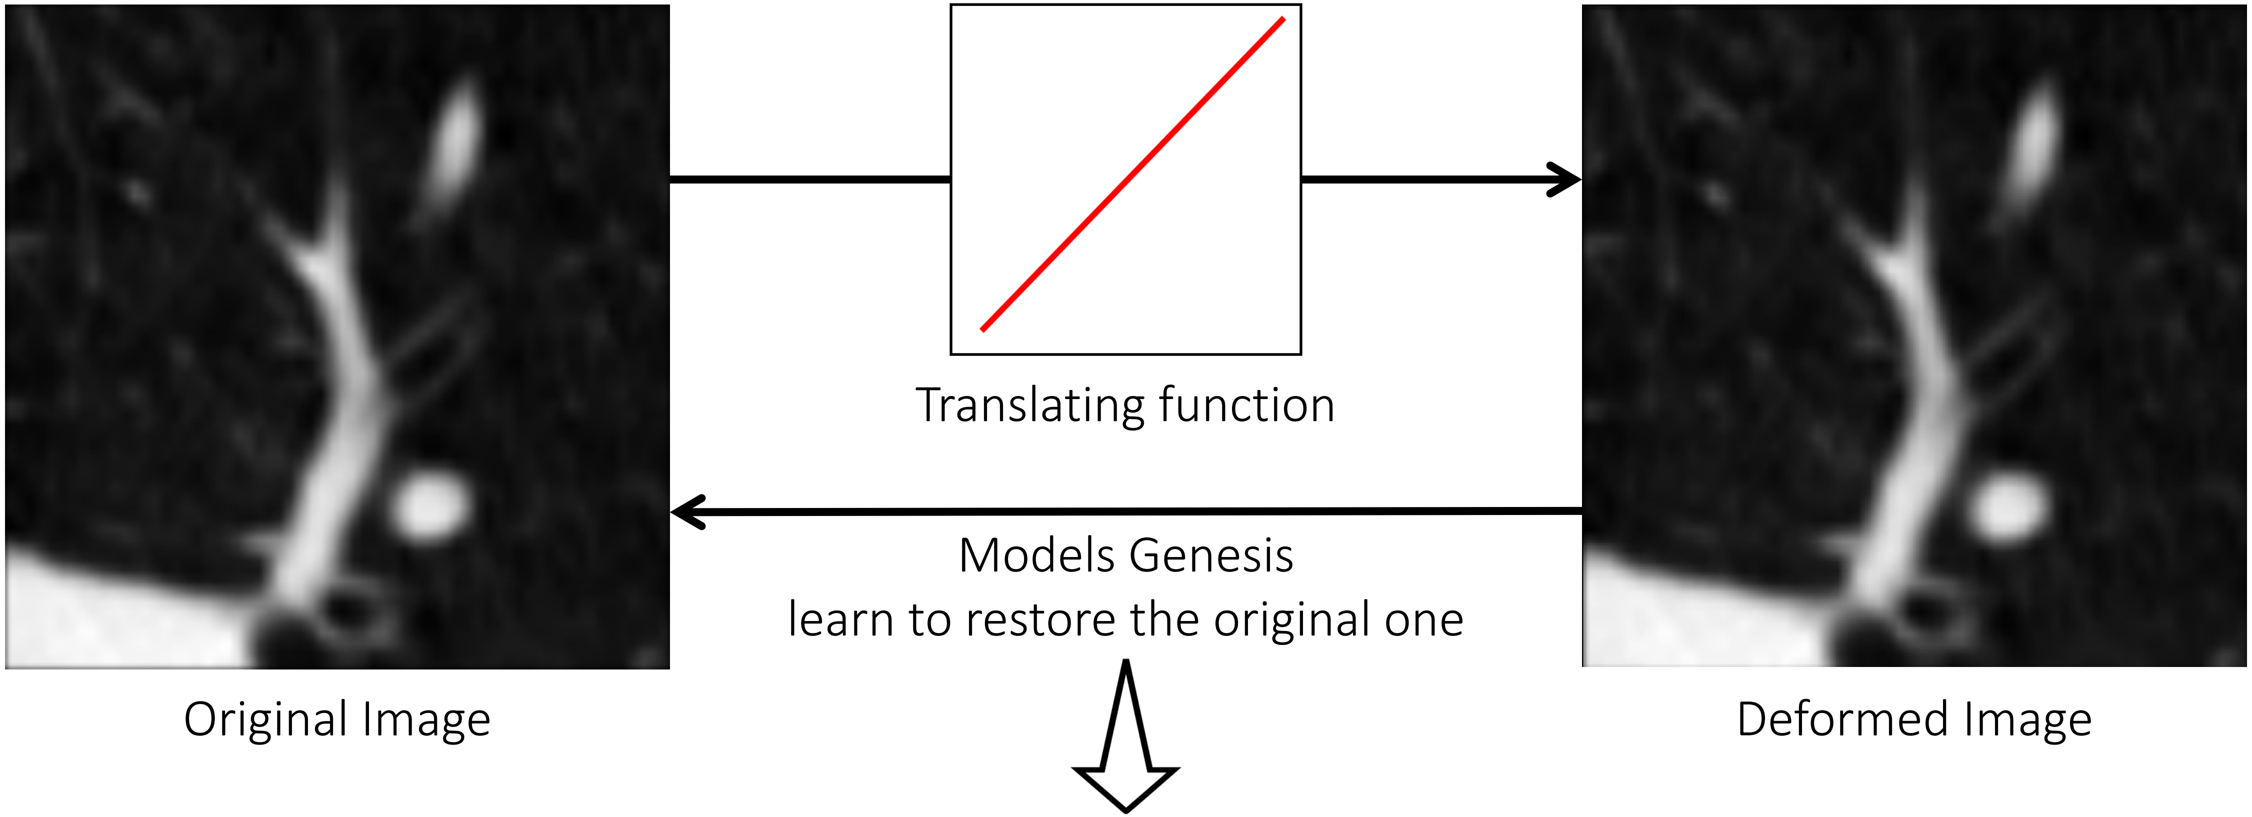

Learning organ appearance and intensity distribution

I. Non-linear transformation

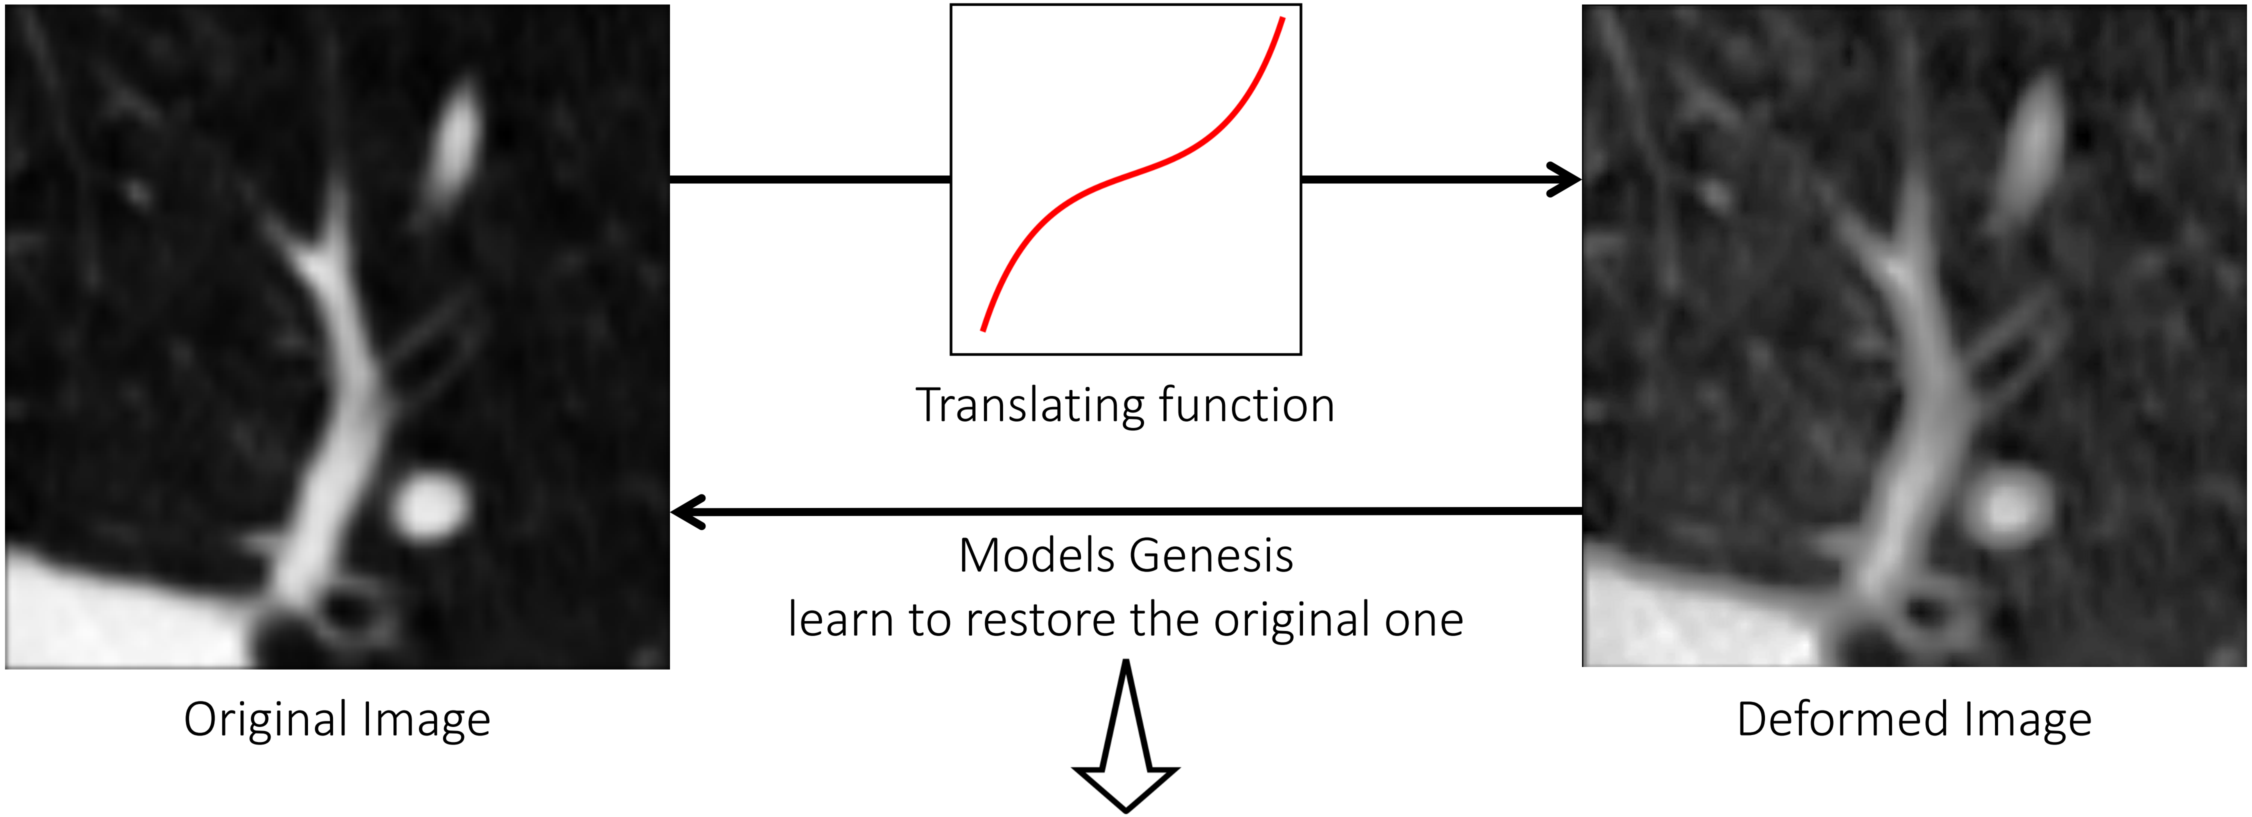

Learning organ appearance and intensity distribution

I. Non-linear transformation

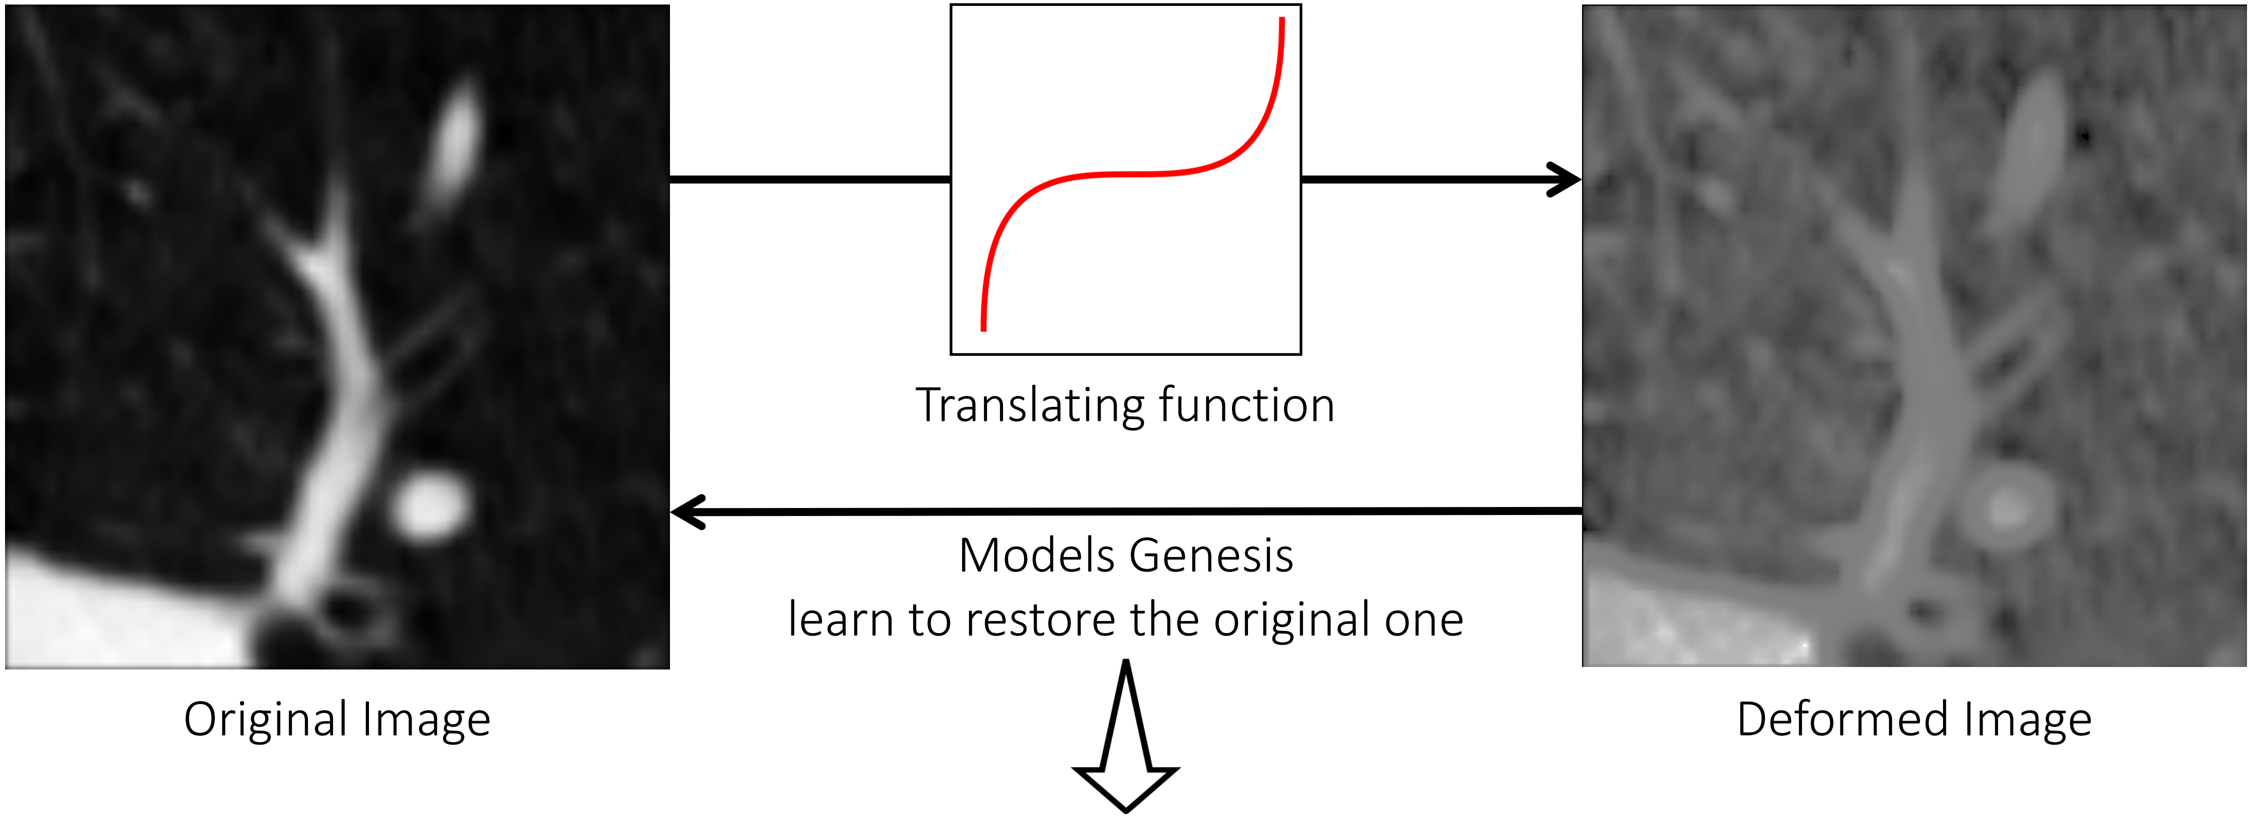

Learning organ appearance and intensity distribution

I. Non-linear transformation

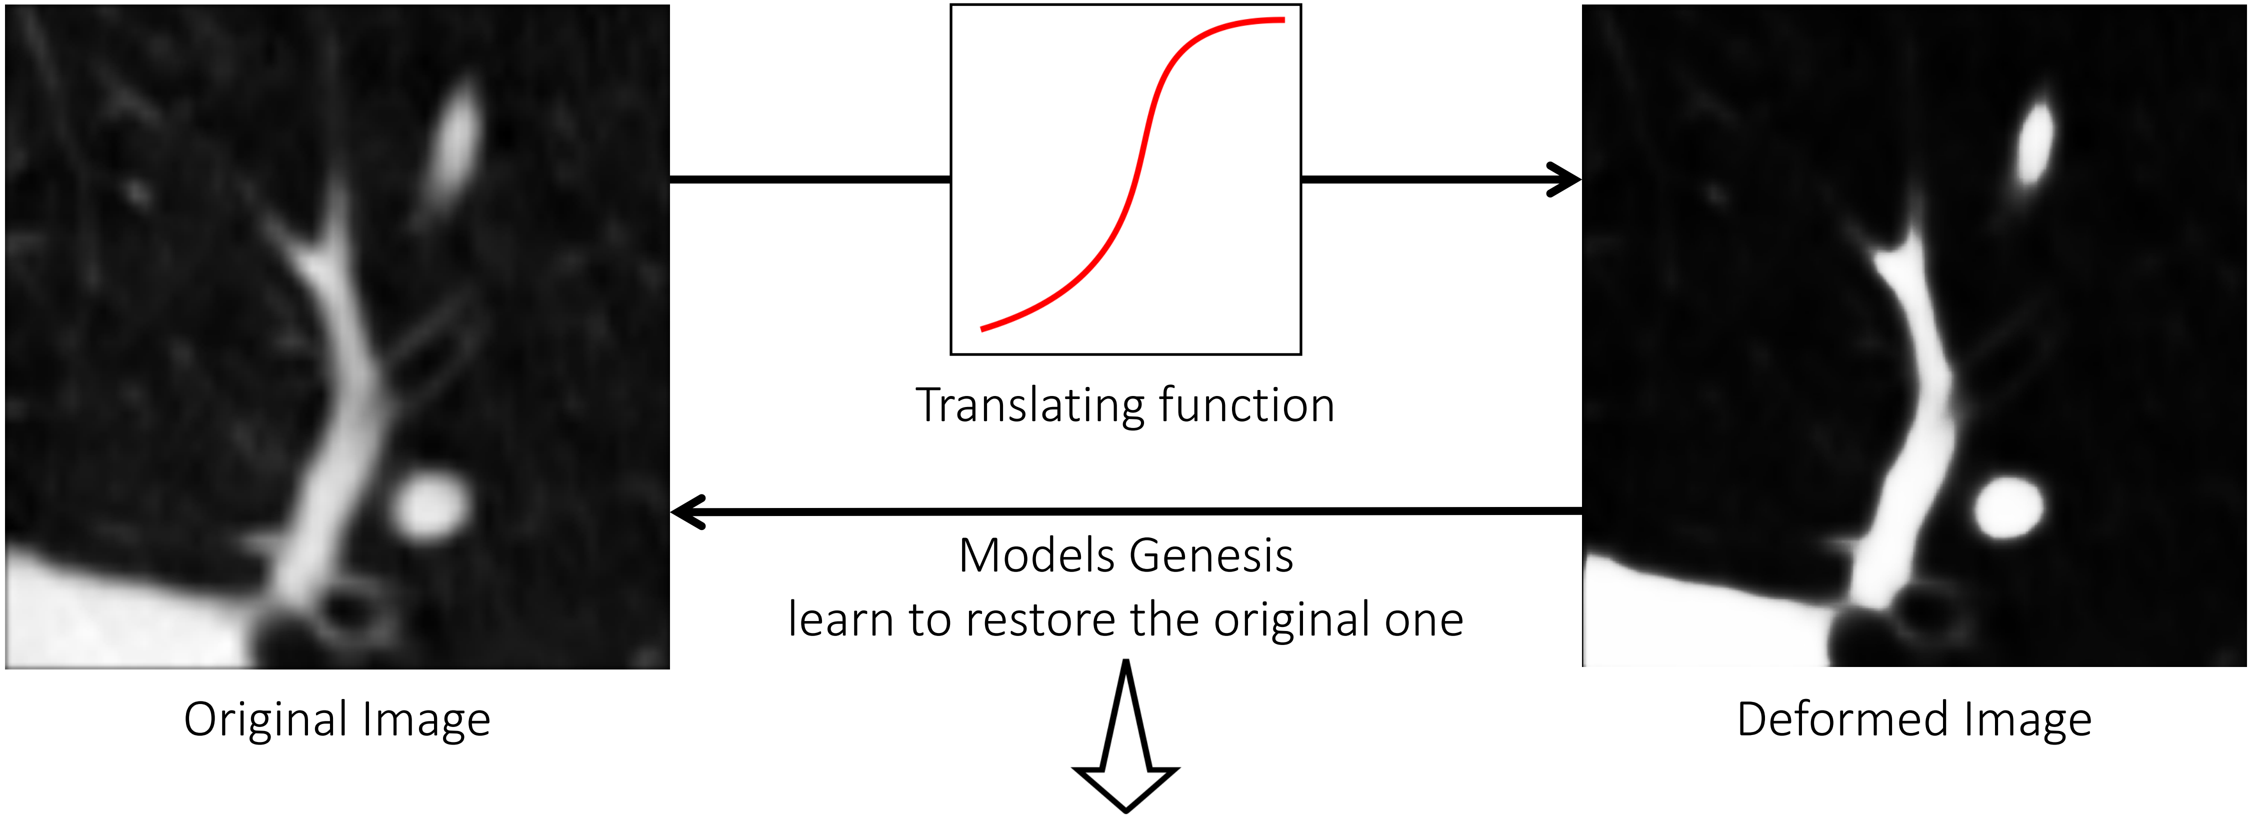

Learning organ appearance and intensity distribution

I. Non-linear transformation

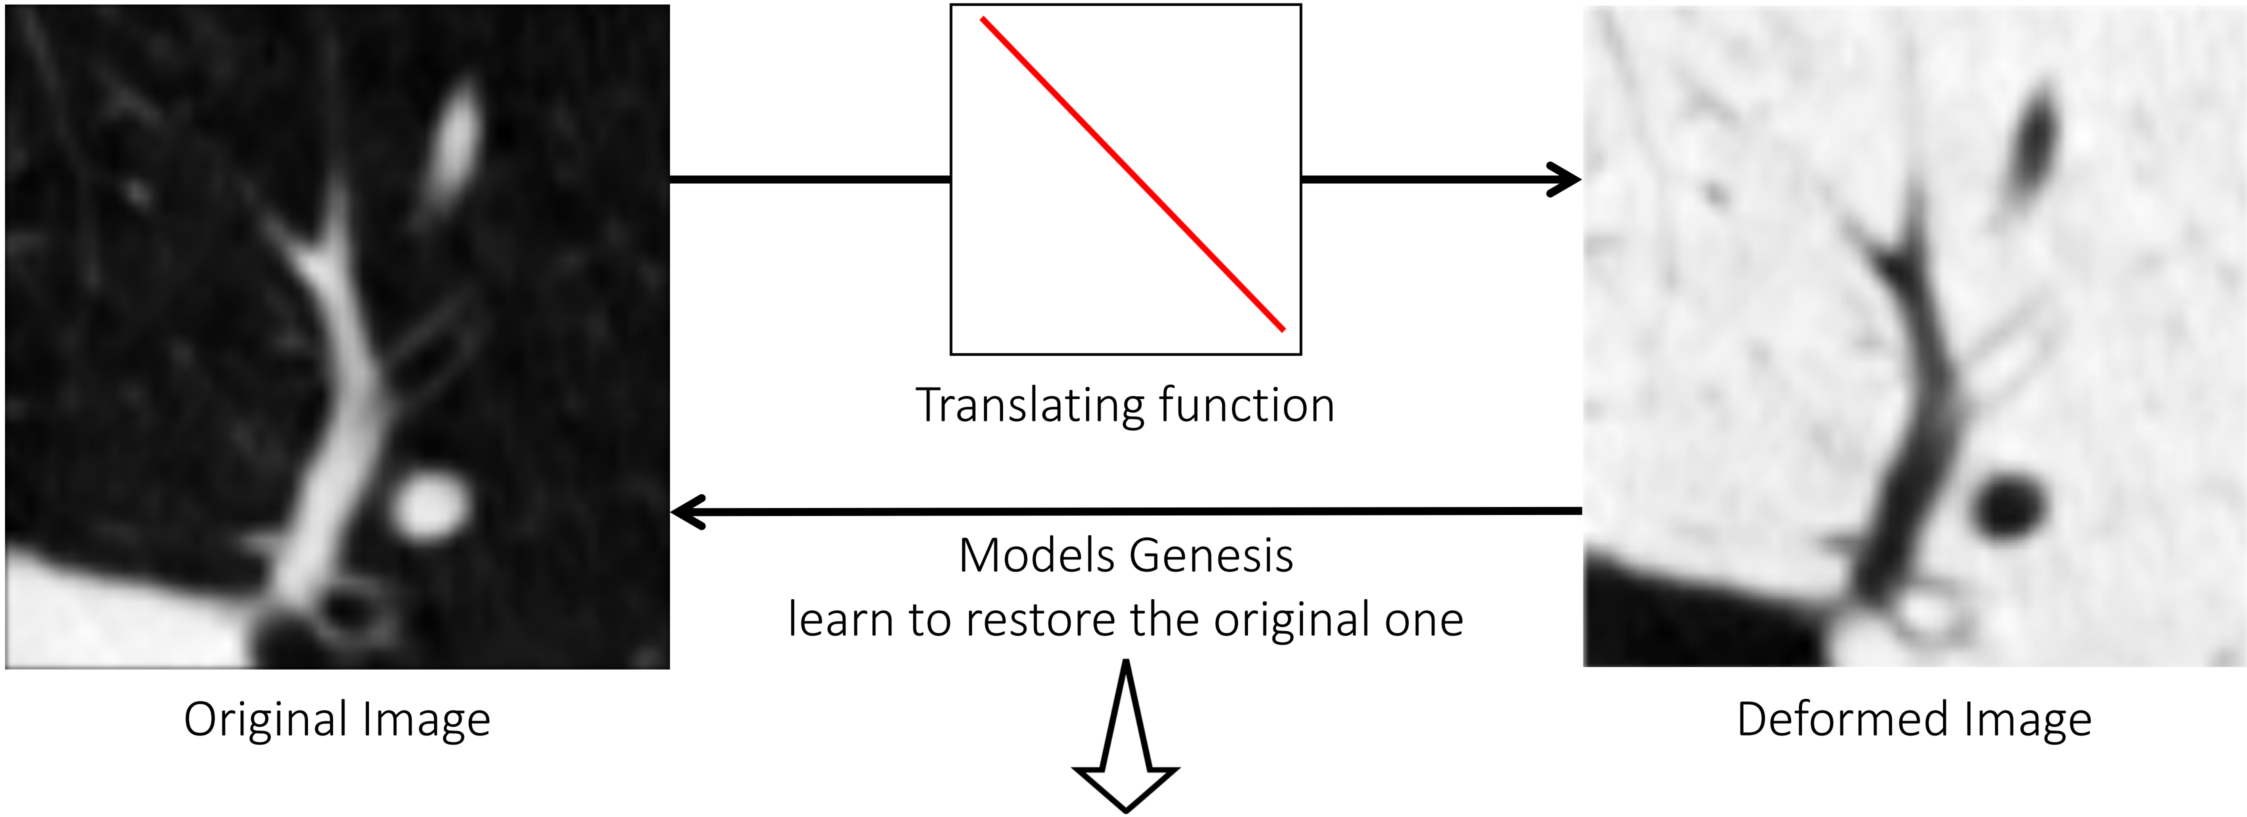

Learning organ appearance and intensity distribution

I. Non-linear transformation

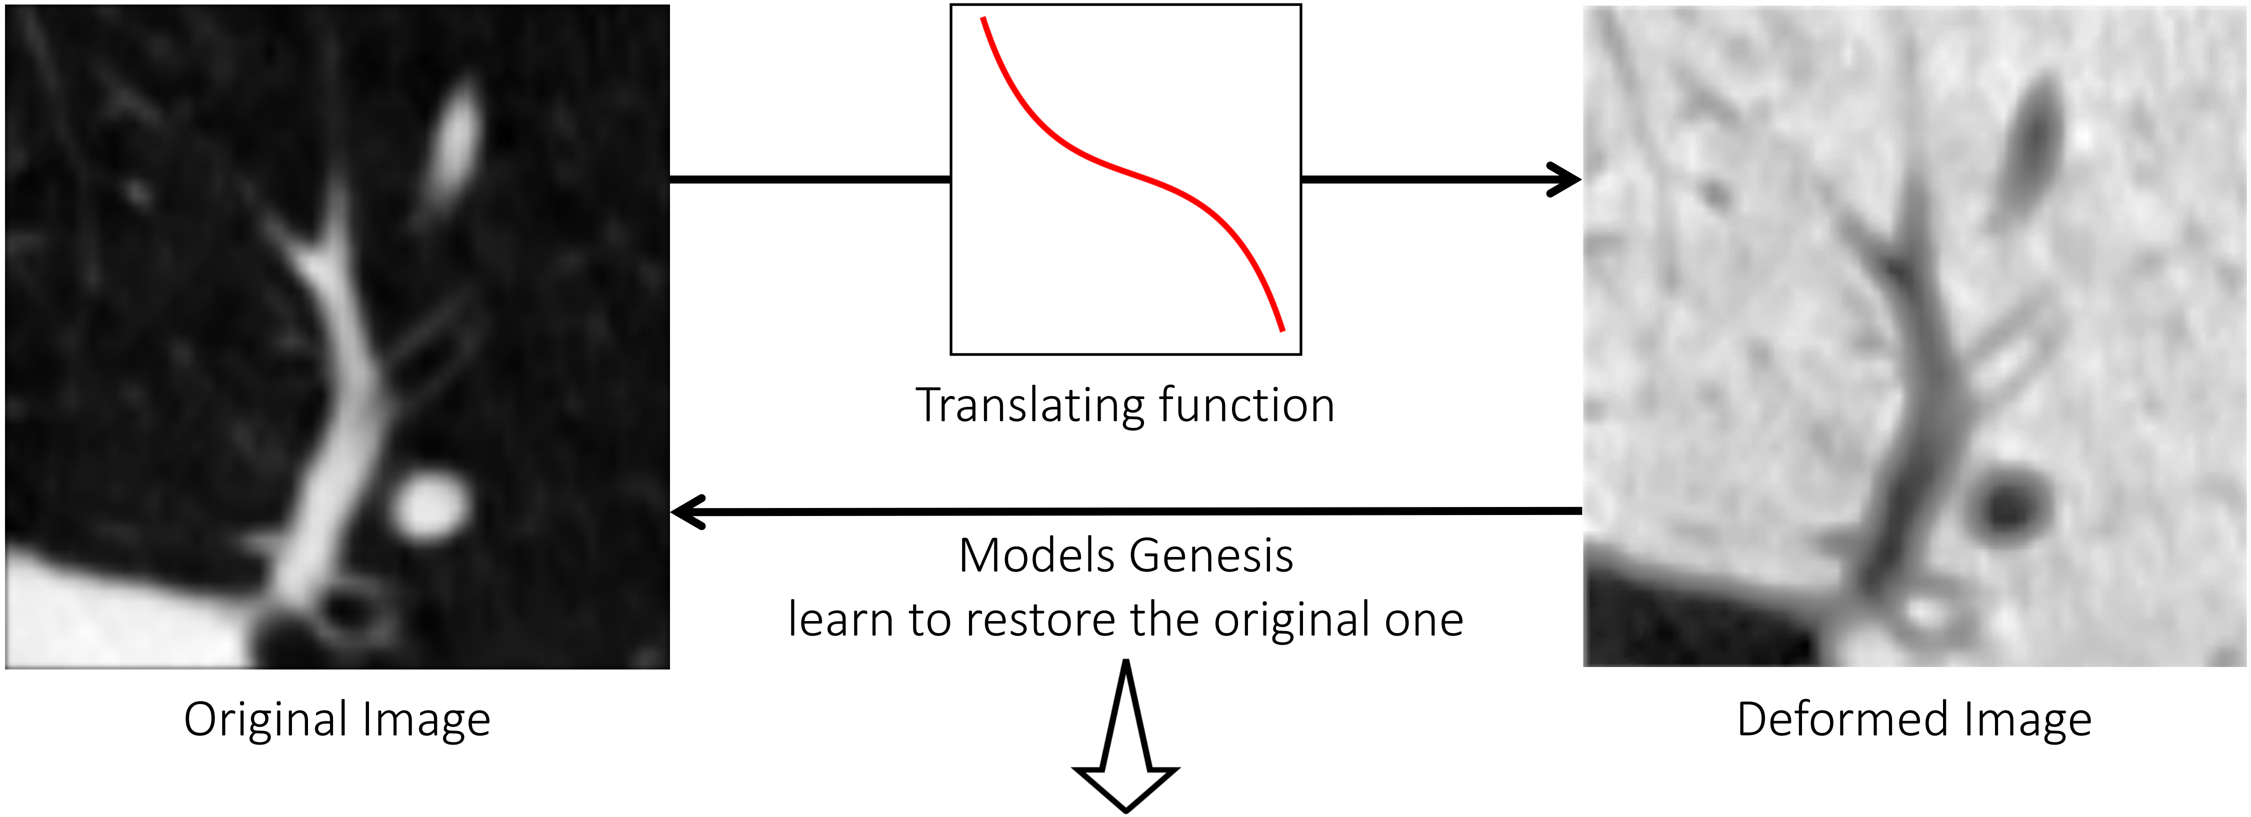

Learning organ appearance and intensity distribution

I. Non-linear transformation

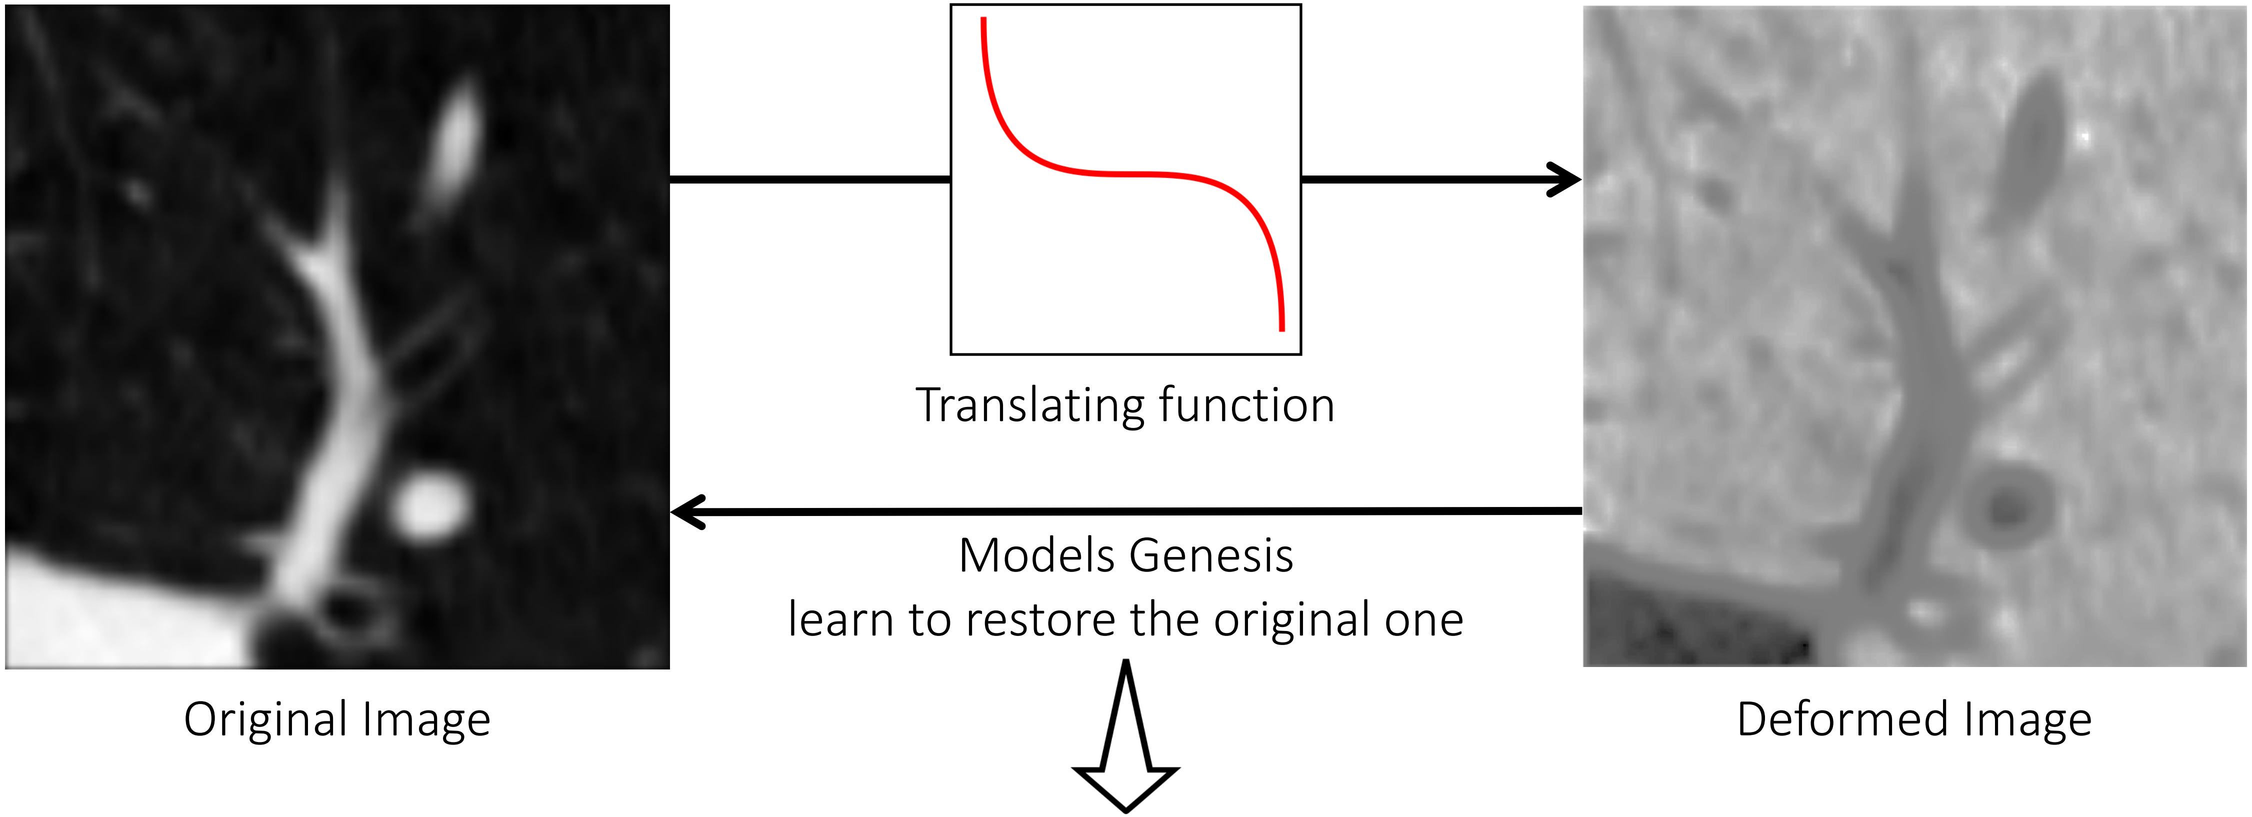

Learning organ appearance and intensity distribution

I. Non-linear transformation

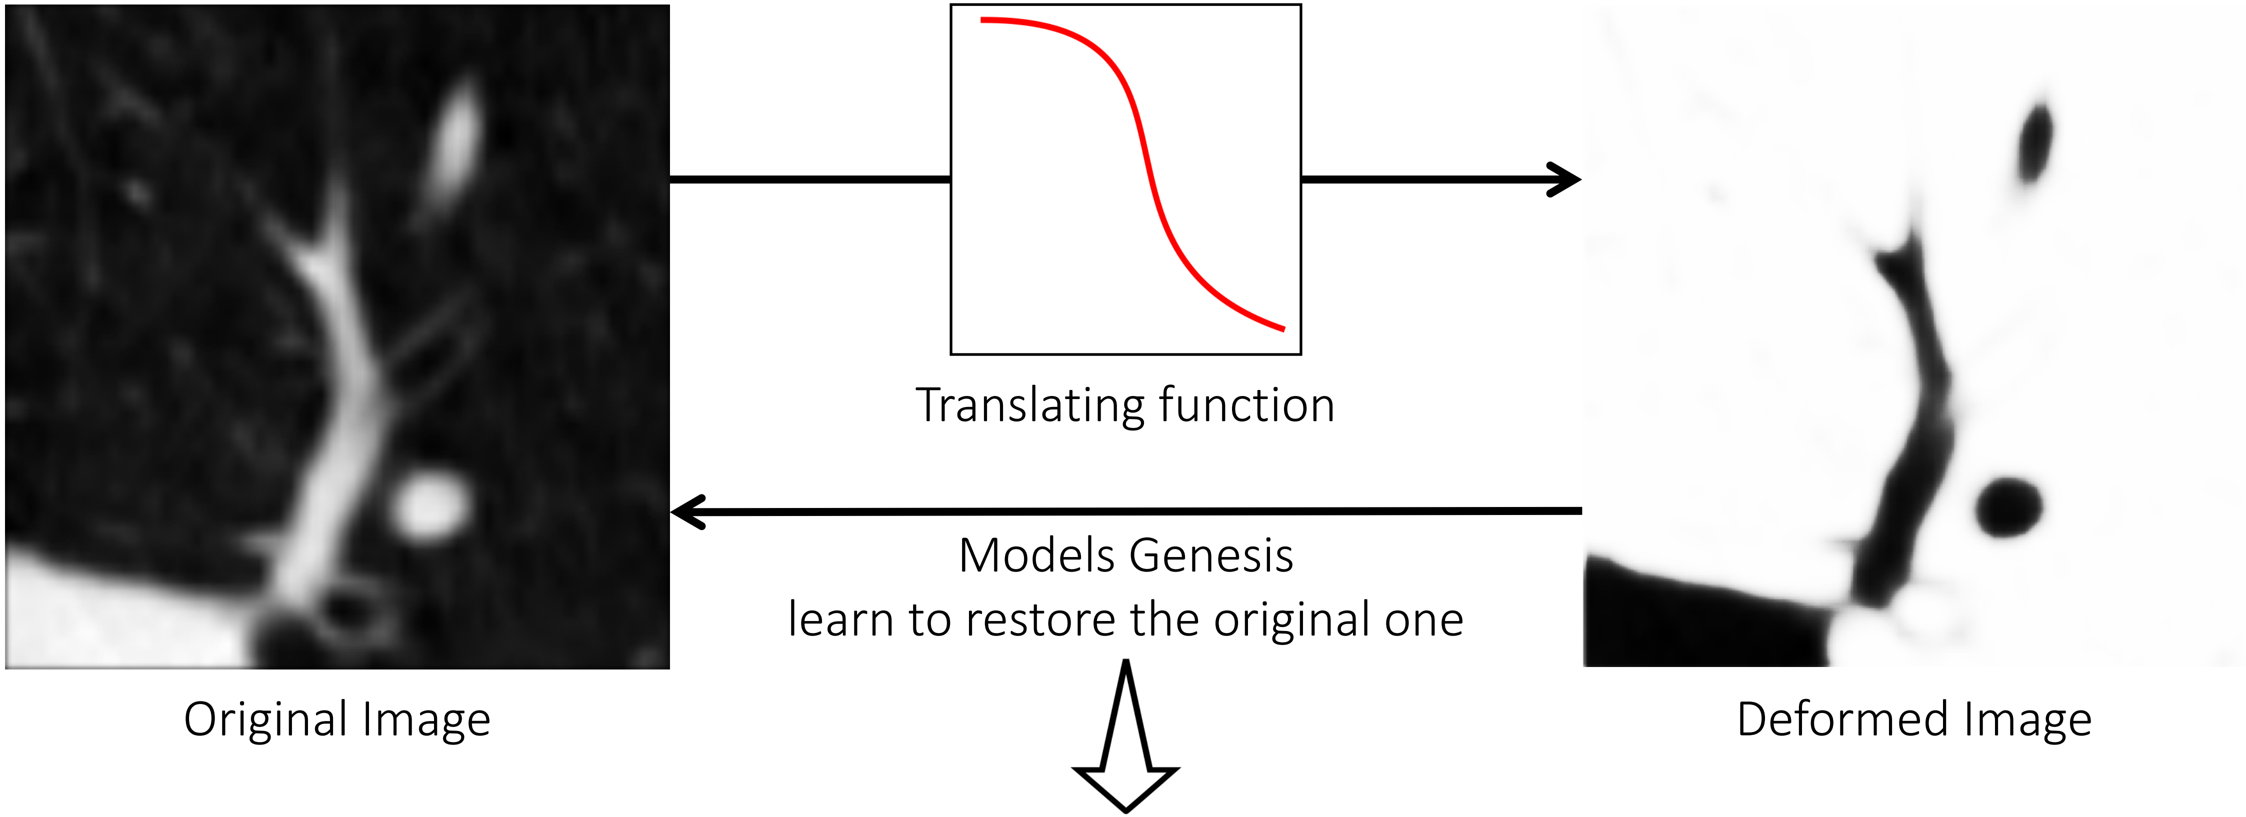

Learning organ appearance and intensity distribution

II. Local pixel shuffling

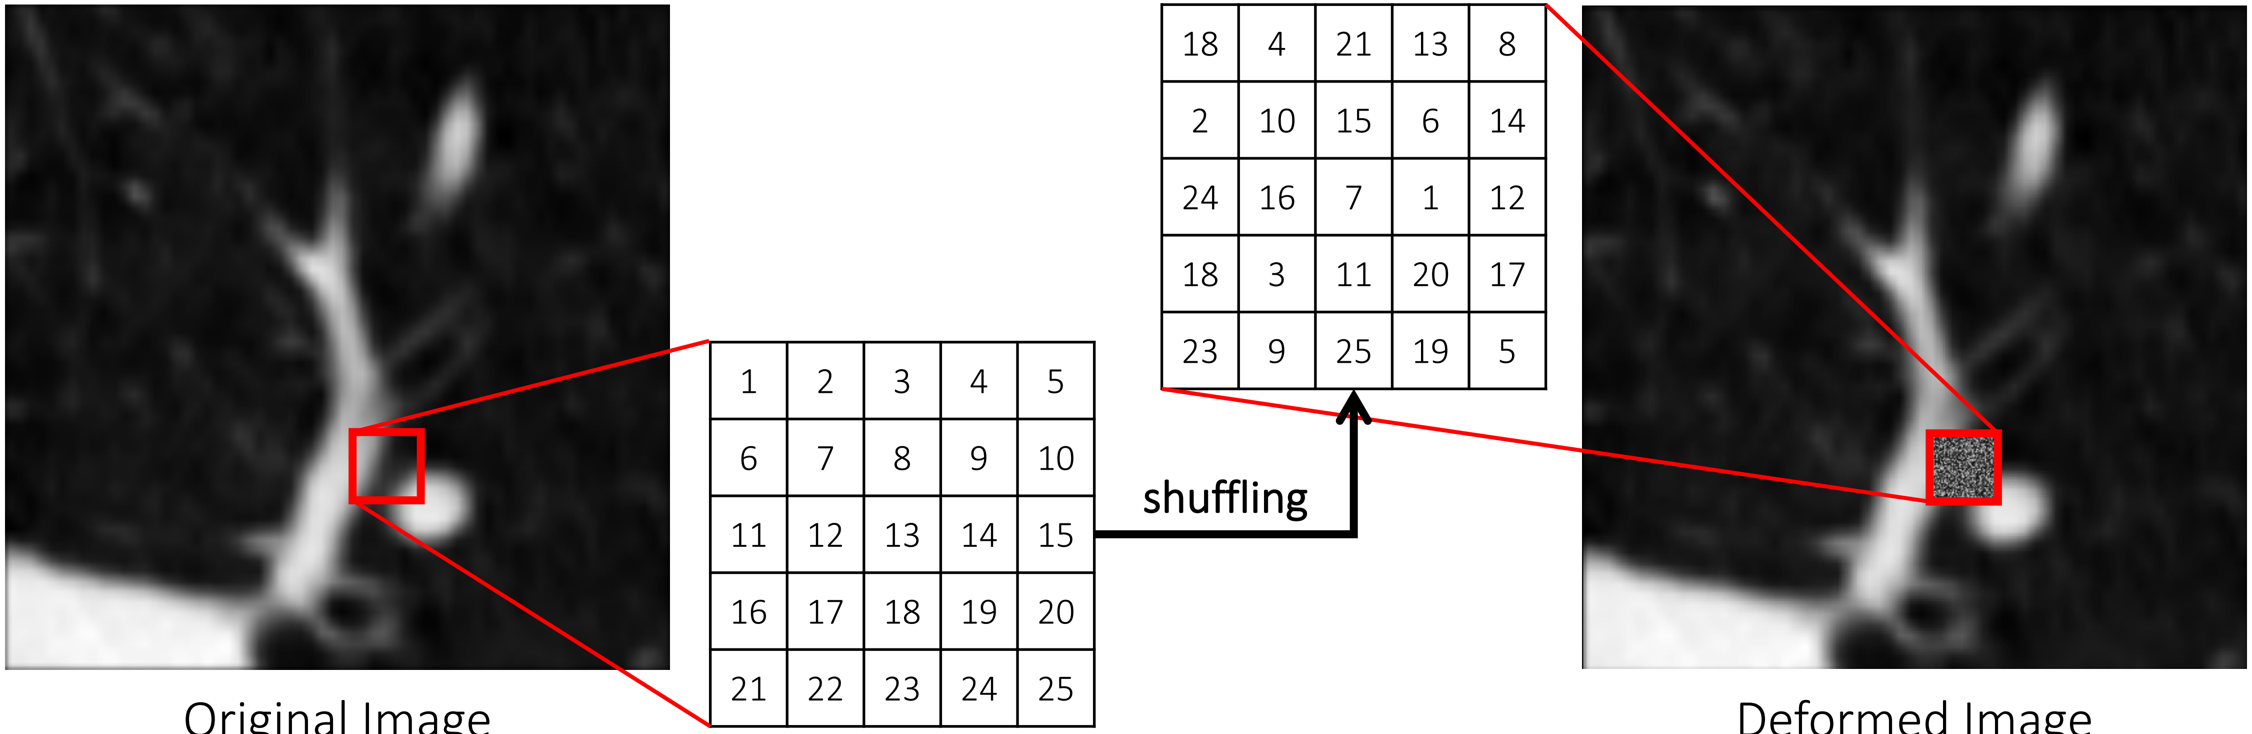

Original Image Deformed Image

II. Local pixel shuffling

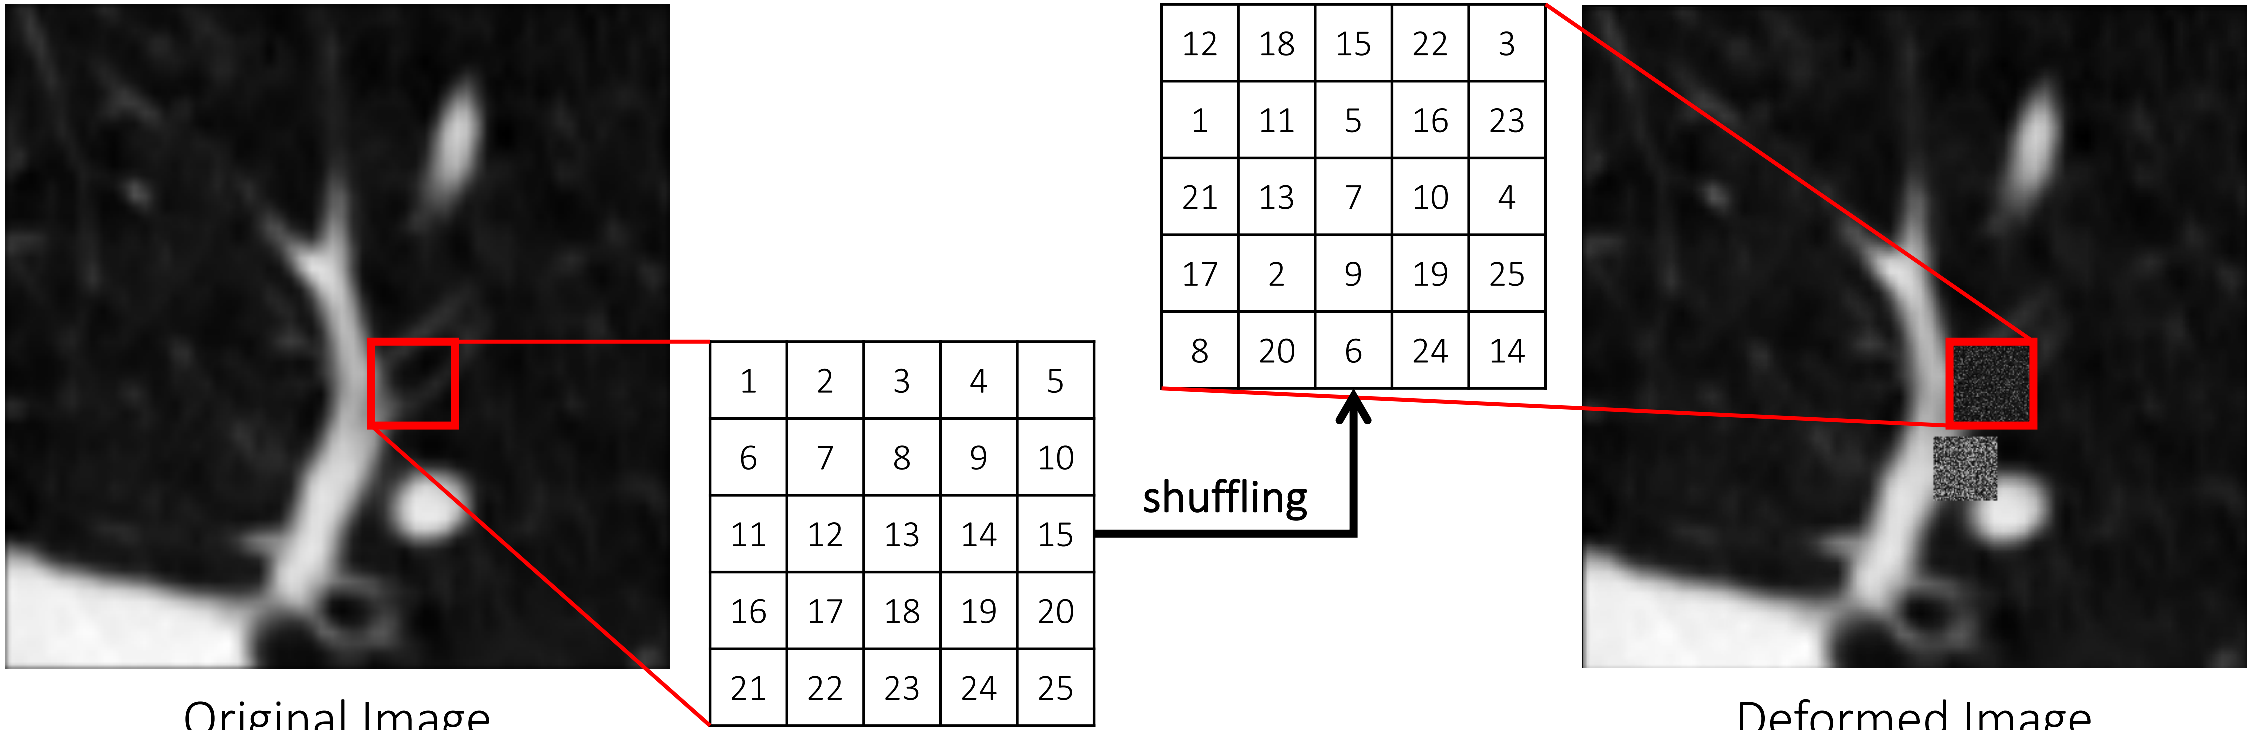

Original Image

Deformed Image

II. Local pixel shuffling

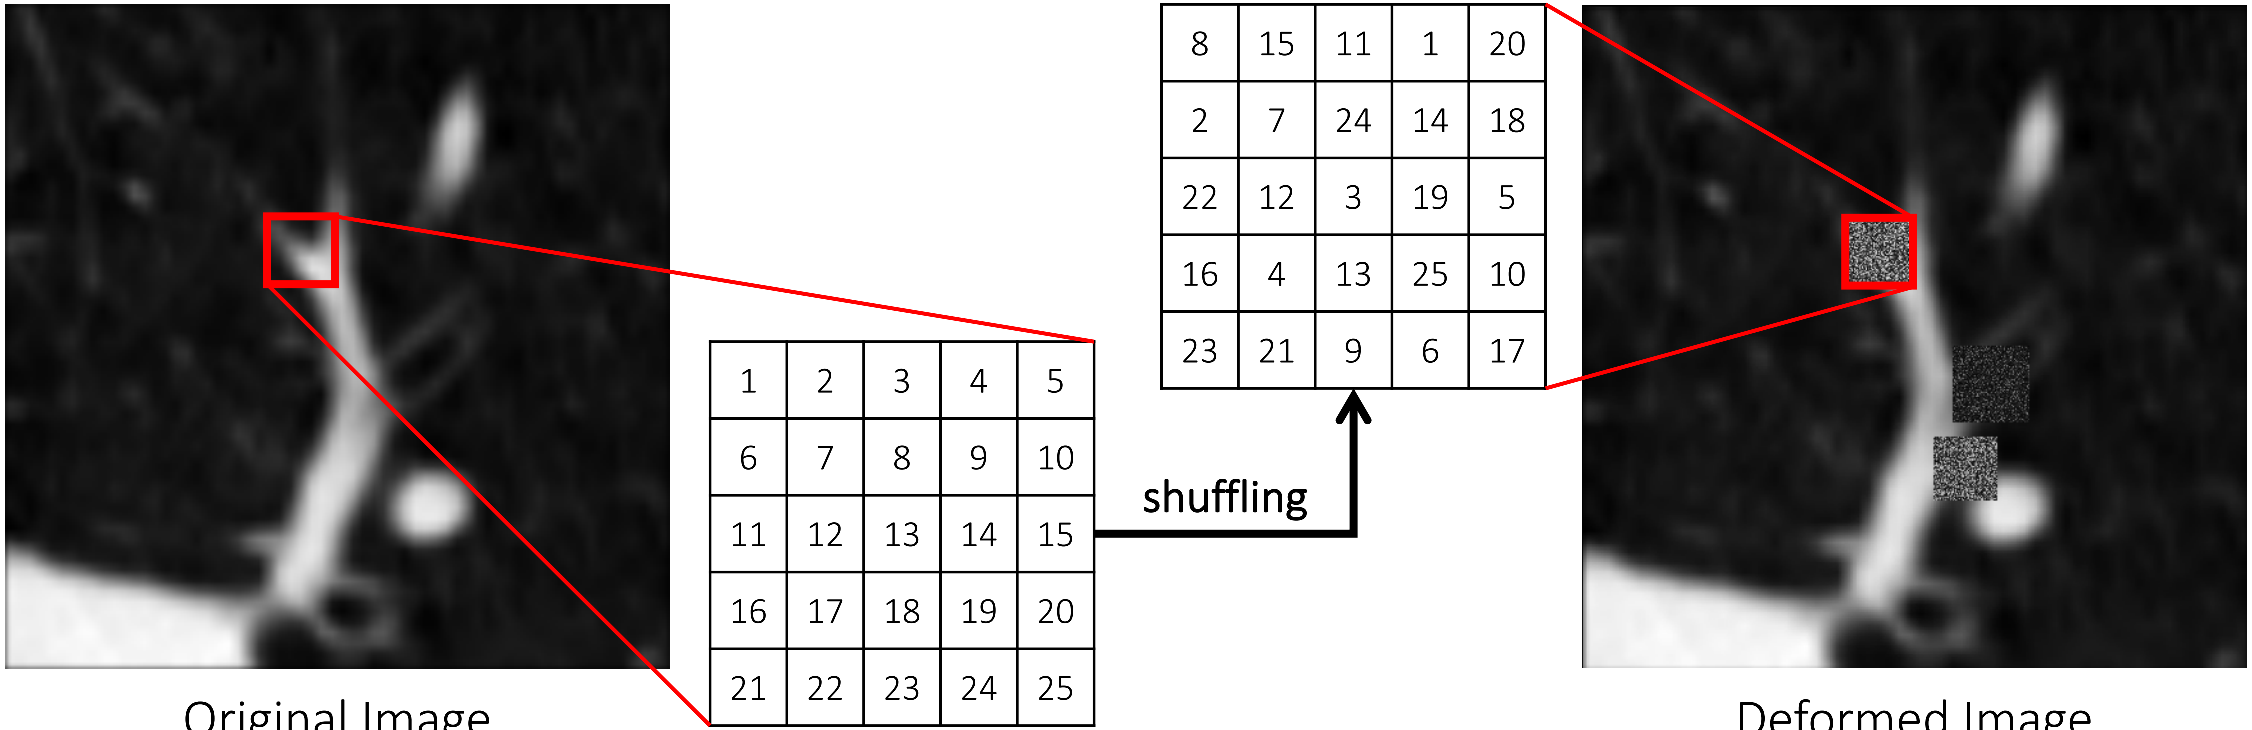

Original Image

Deformed Image

II. Local pixel shuffling

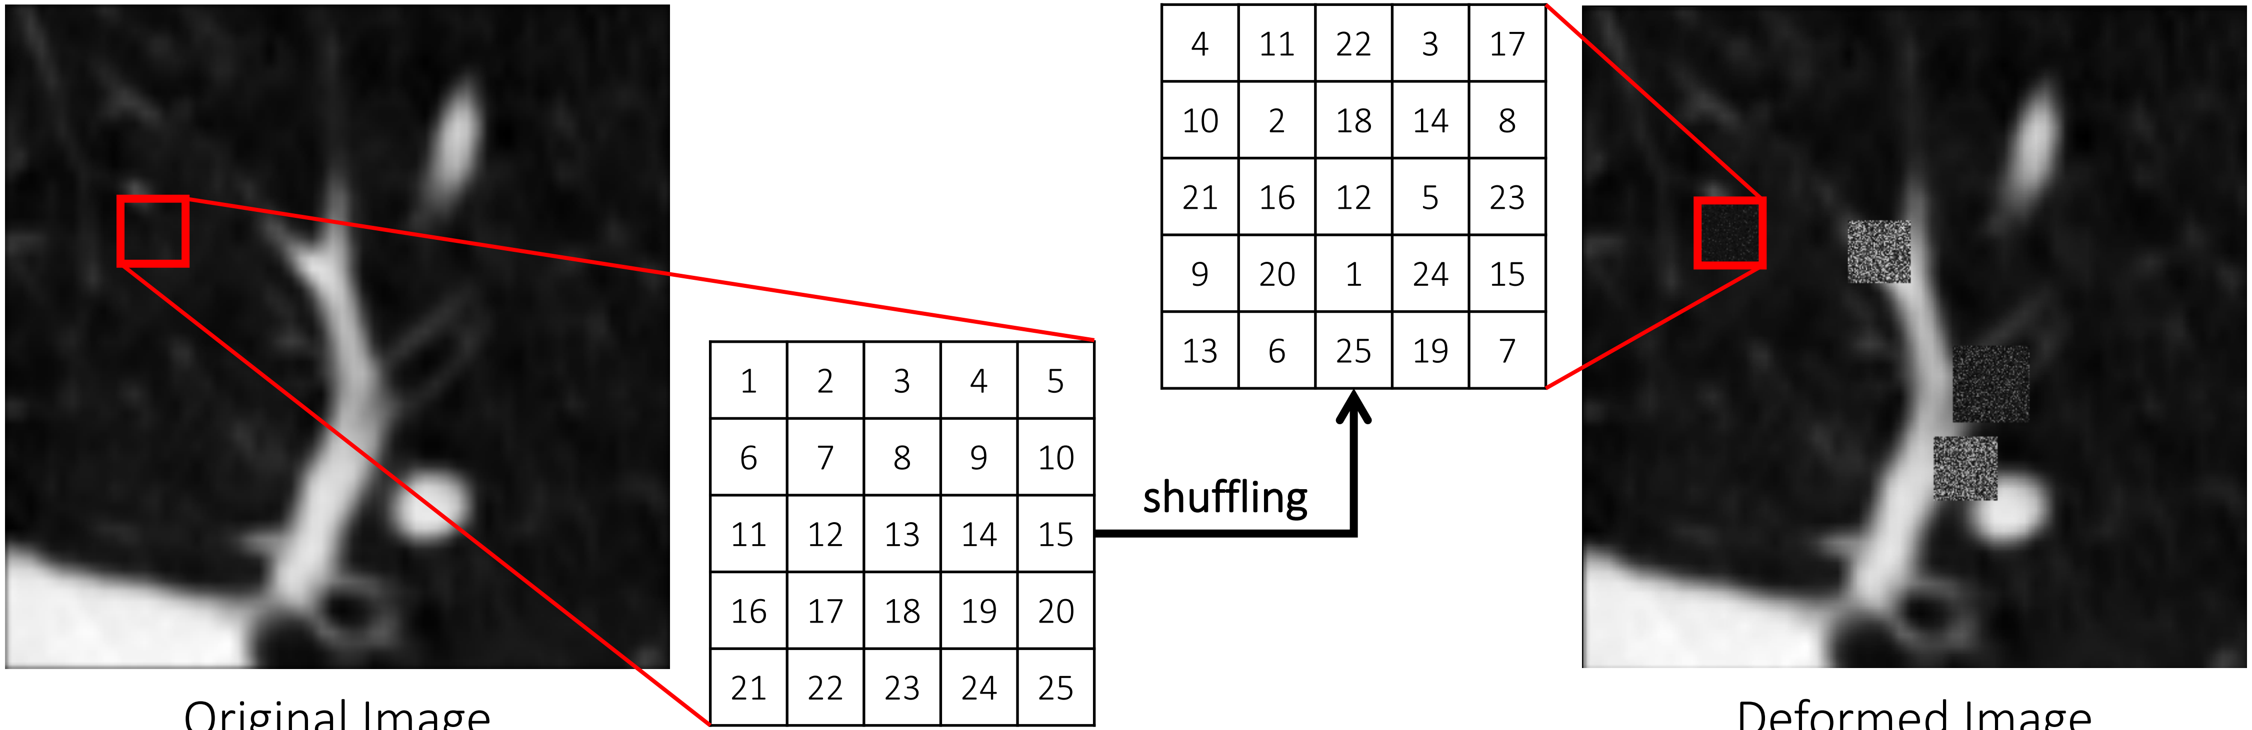

Original Image

Deformed Image

II. Local pixel shuffling

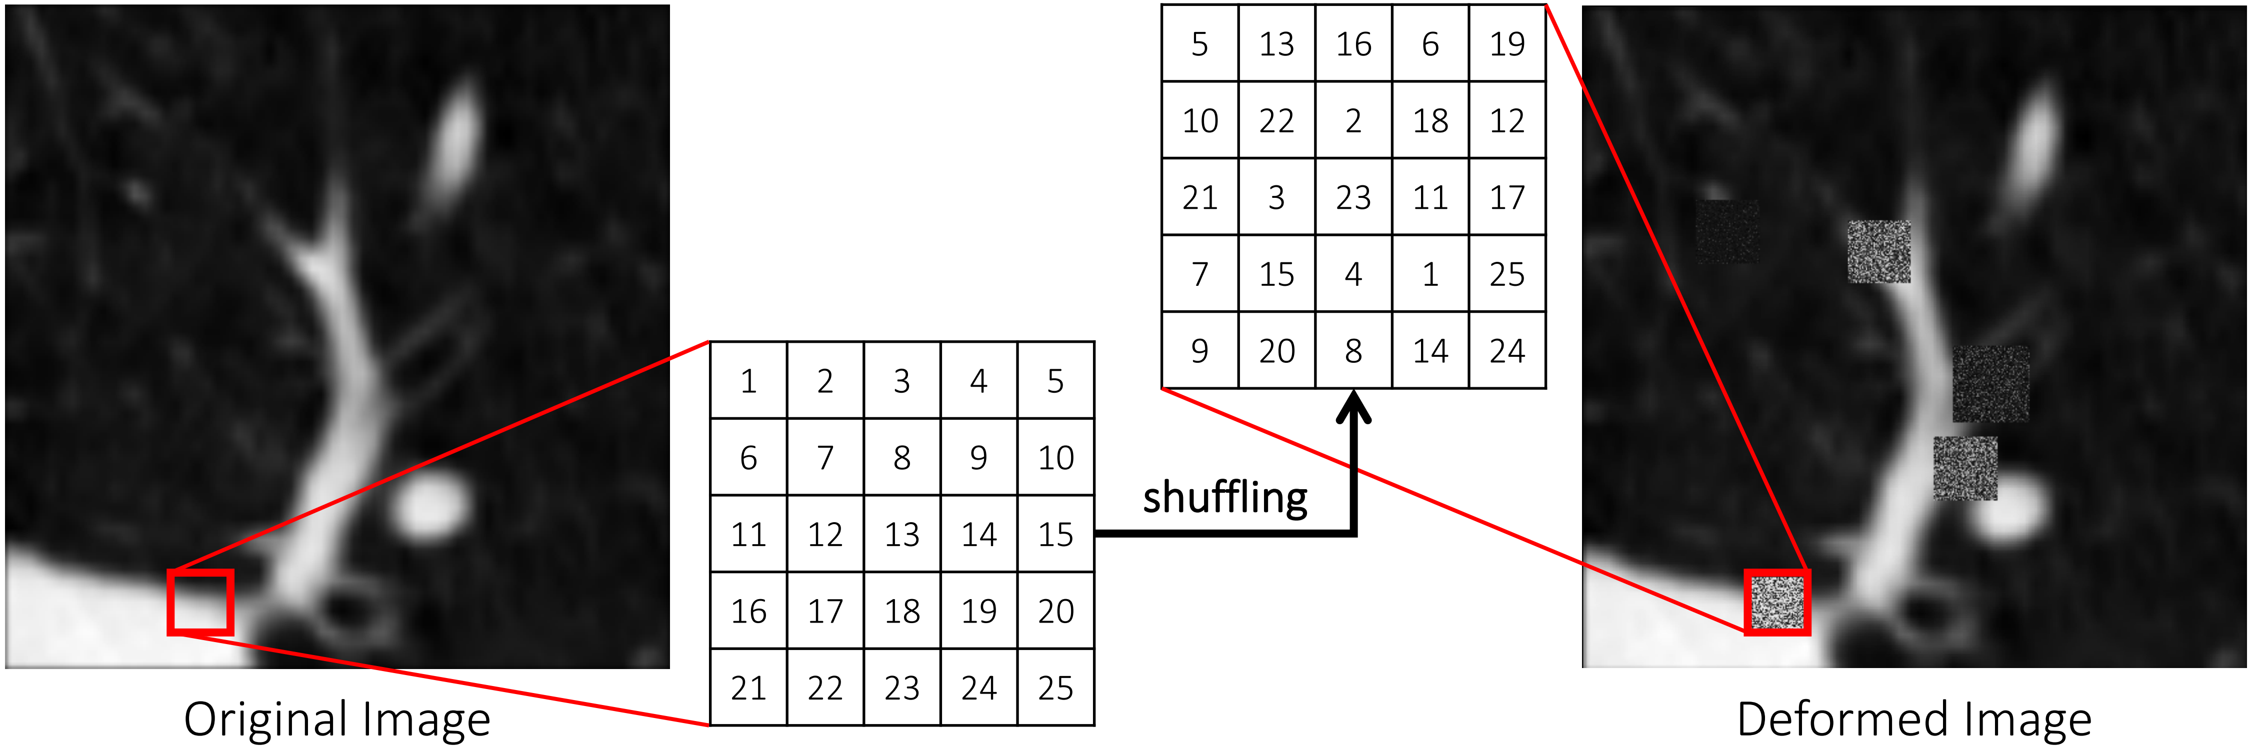

II. Local pixel shuffling

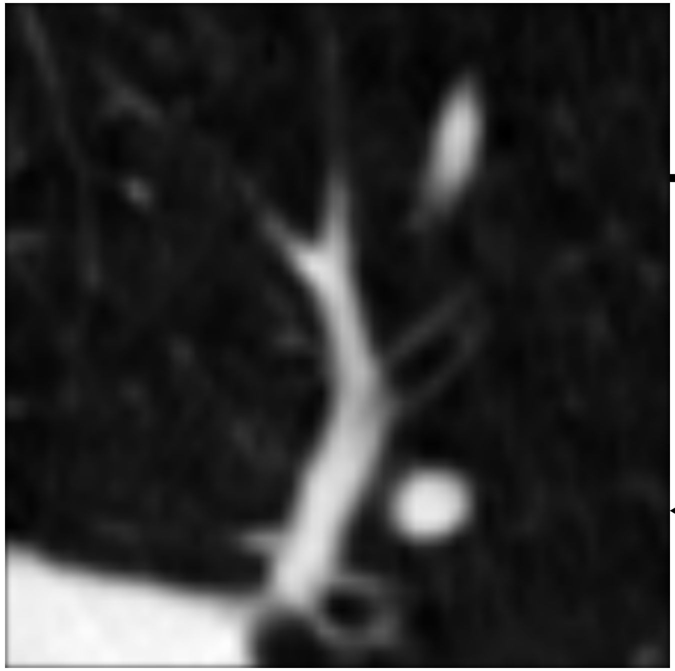

Shuffling

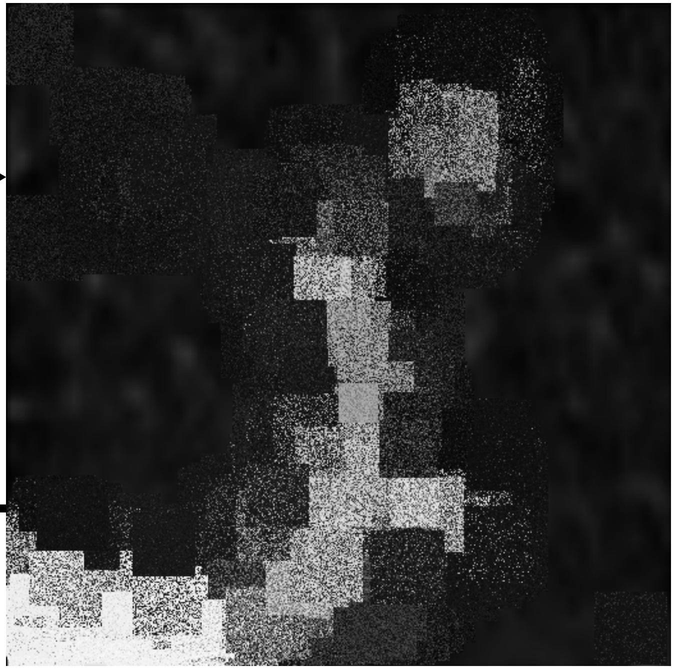

Models Genesis
learn to restore the original one

Original Image

Deformed Image

Learning organ texture and local boundaries

III. Out-painting

IV. In-painting

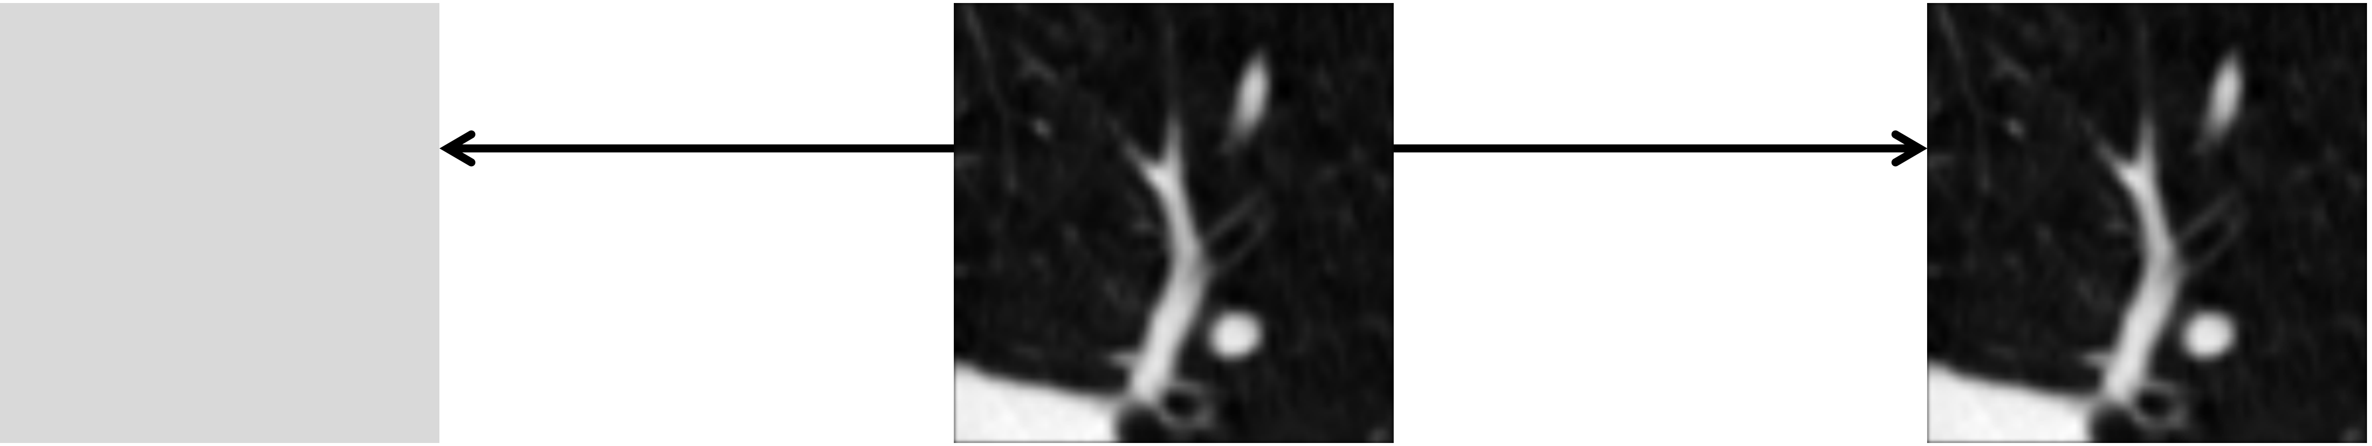

Deformed Image | Original Image | Deformed Image

III. Out-painting

IV. In-painting

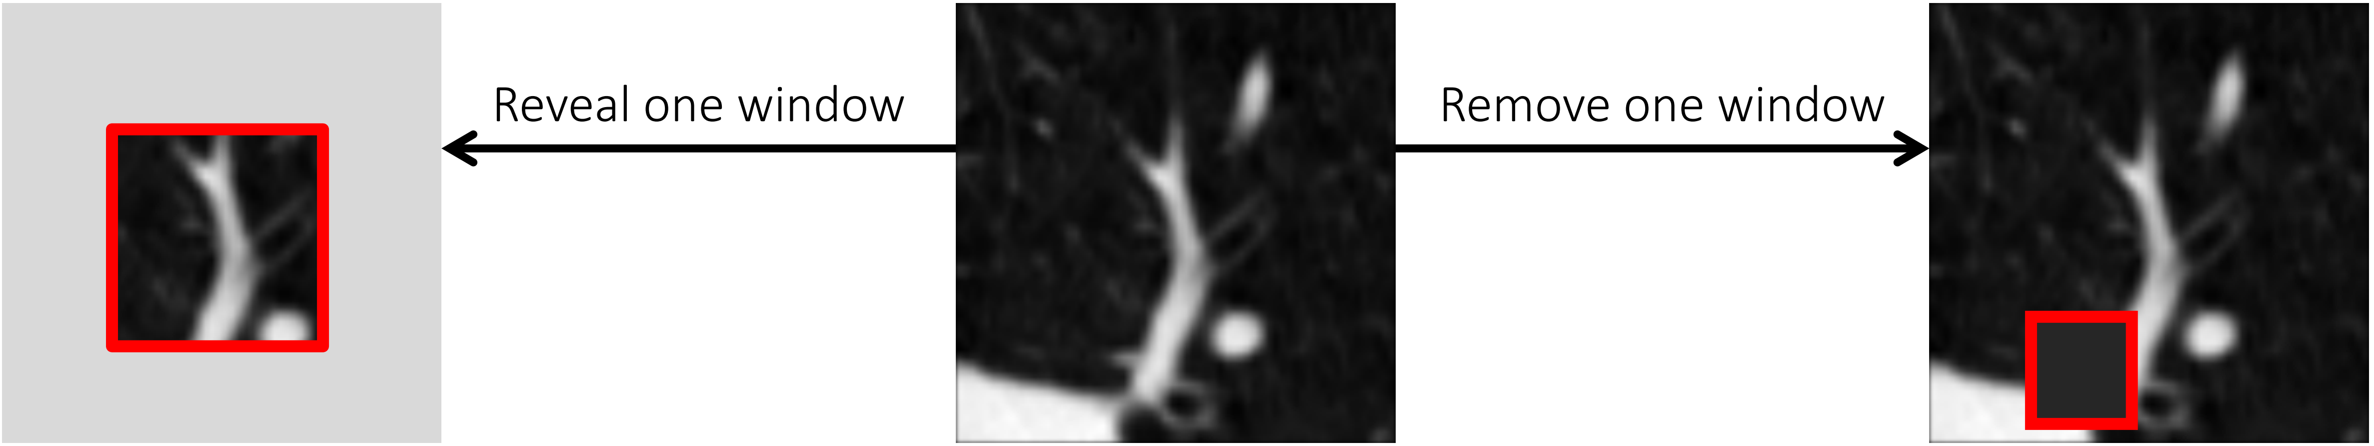

Deformed Image | Original Image | Deformed Image

III. Out-painting

IV. In-painting

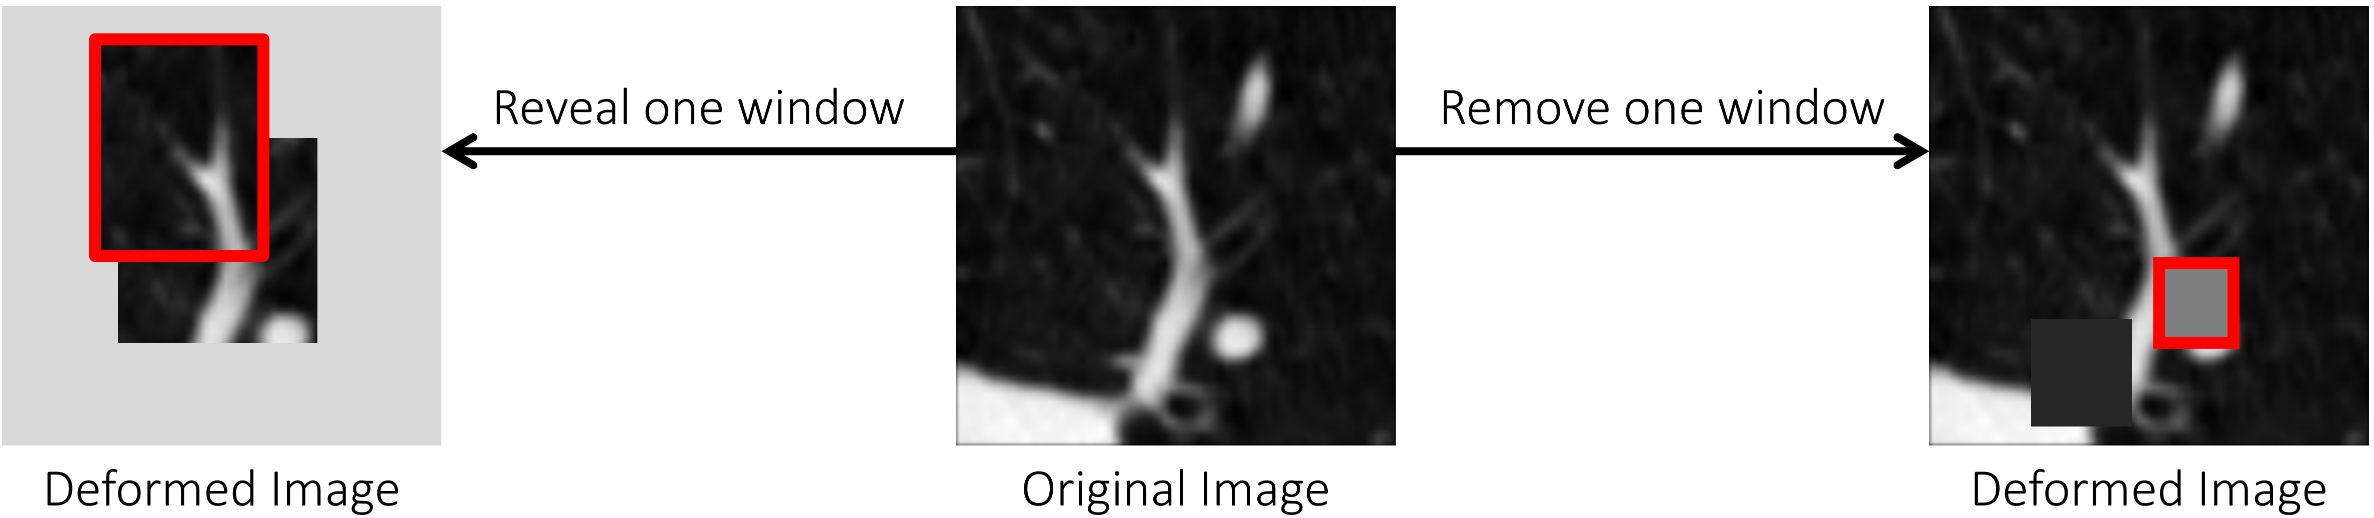

III. Out-painting

IV. In-painting

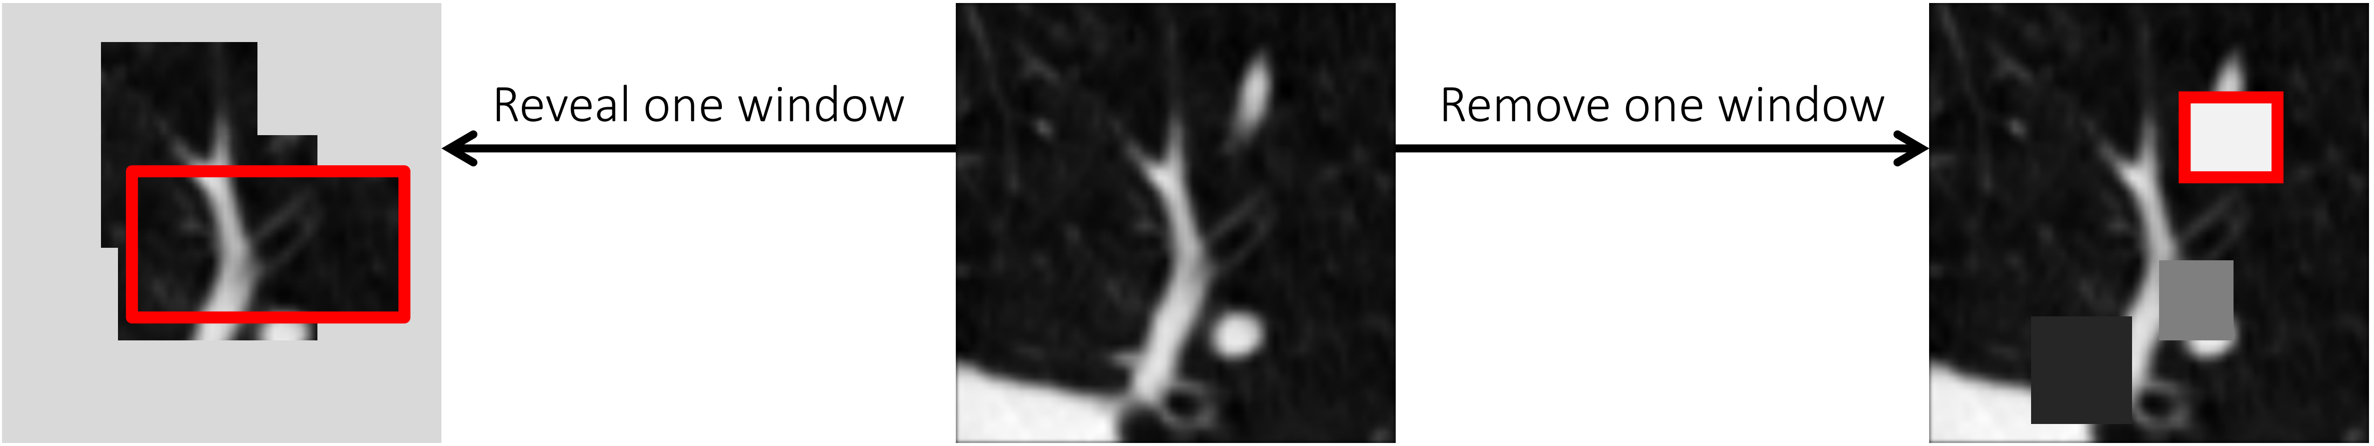

Deformed Image

Original Image

Deformed Image

III. Out-painting

IV. In-painting

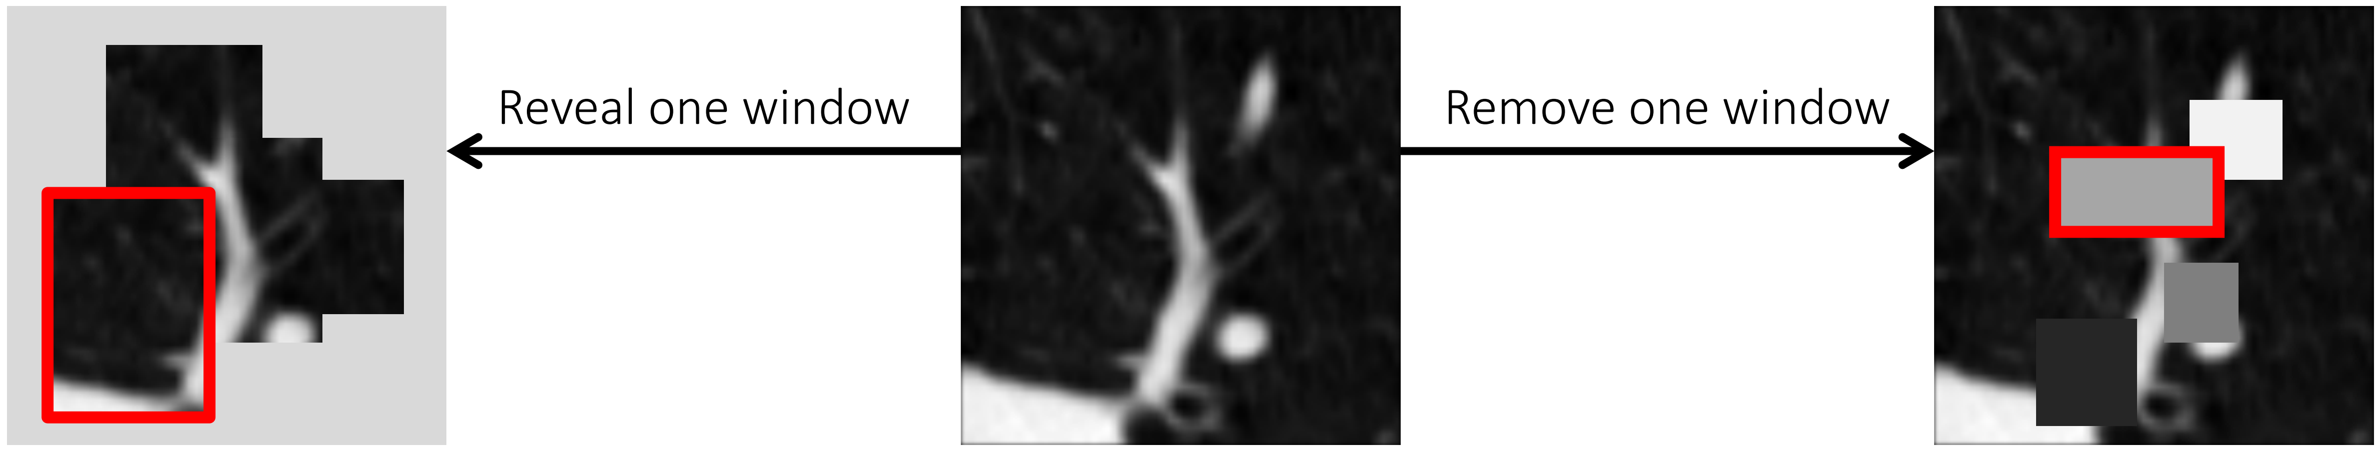

Deformed Image | Original Image | Deformed Image

III. Out-painting

IV. In-painting

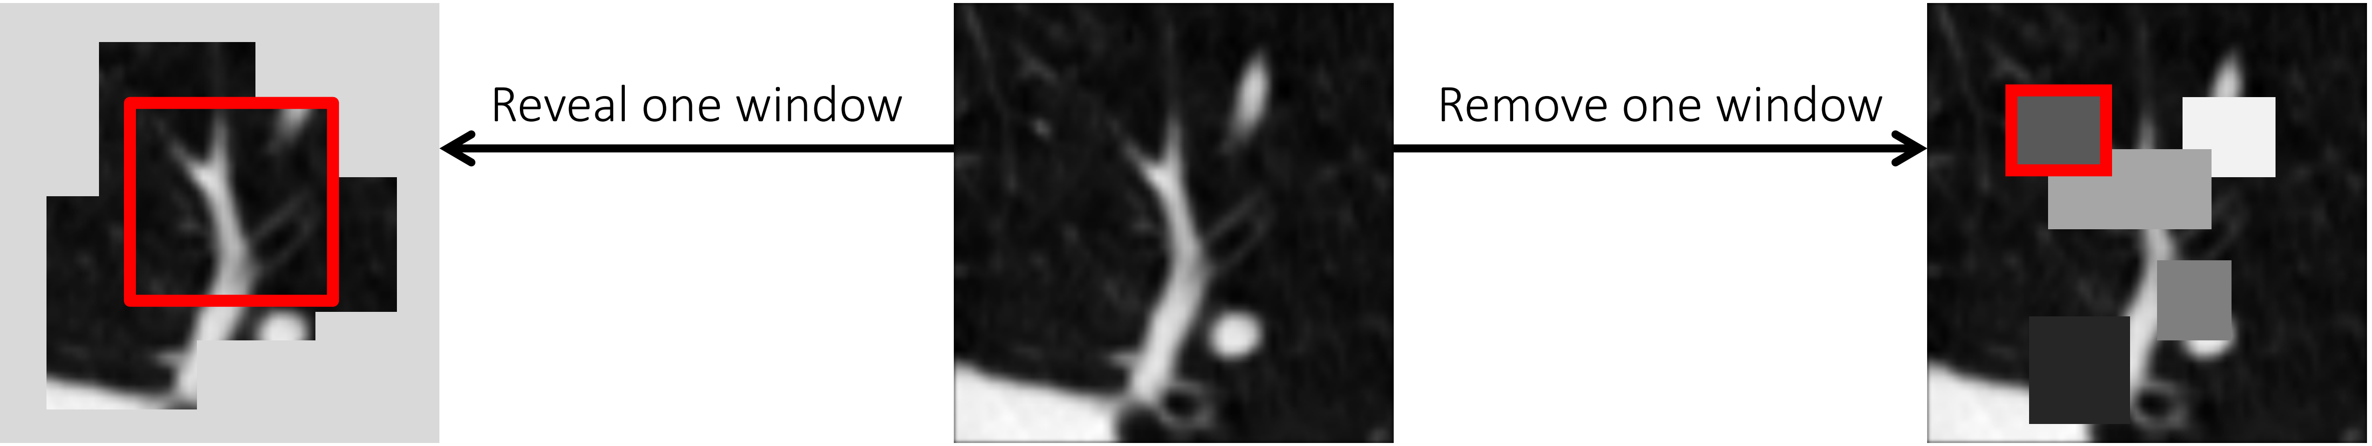

Deformed Image | Original Image | Deformed Image

III. Out-painting

IV. In-painting

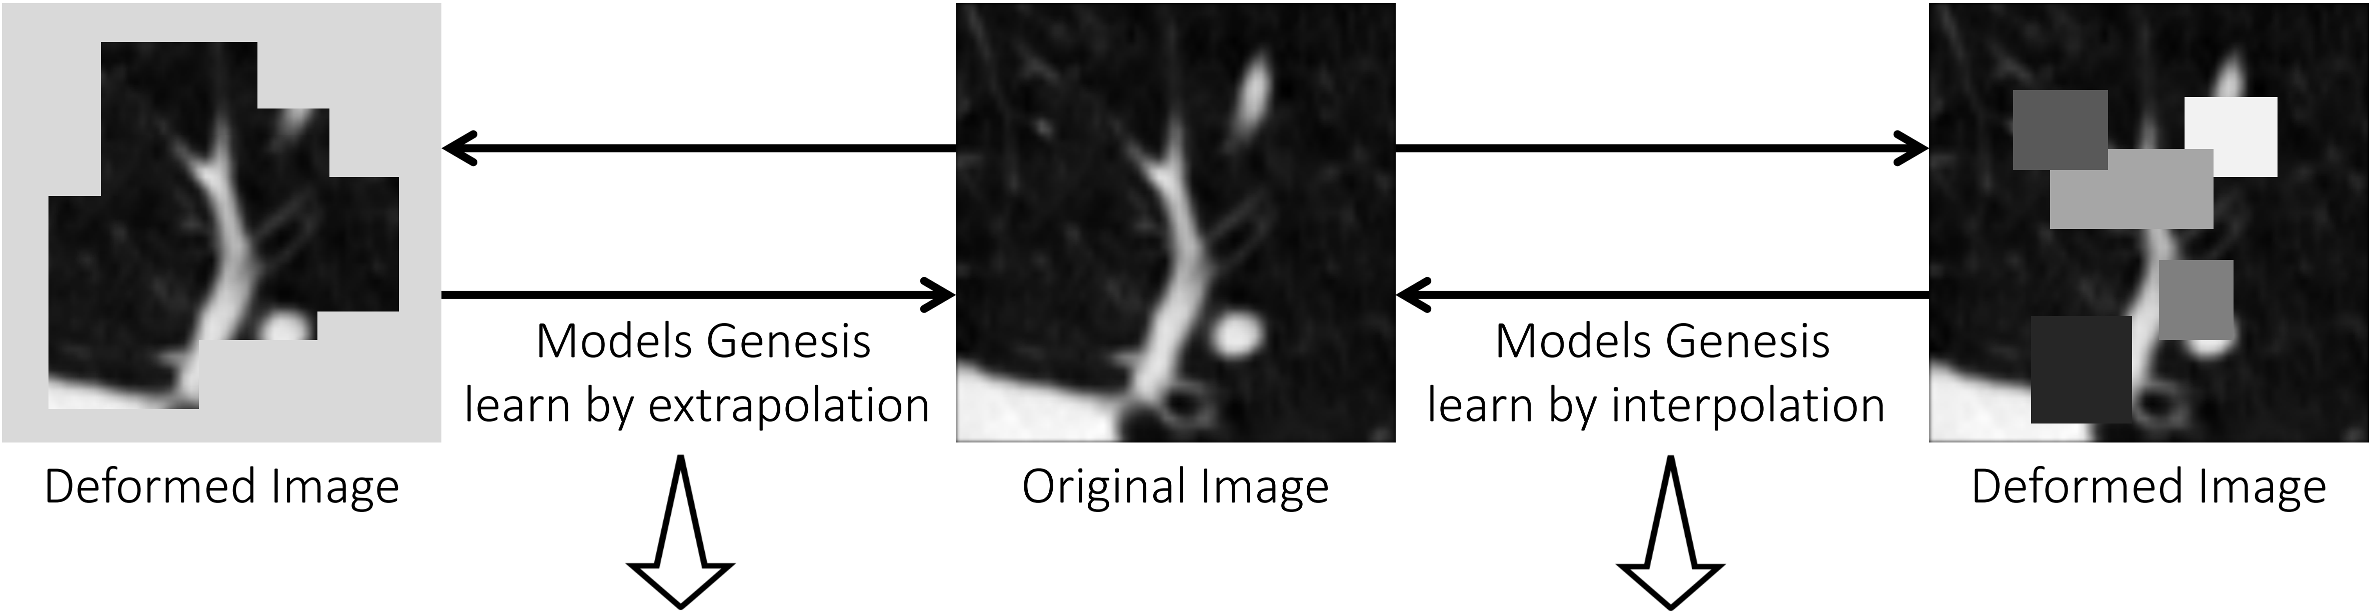

Learning organ spatial layout and global geometry

Learning local continuities

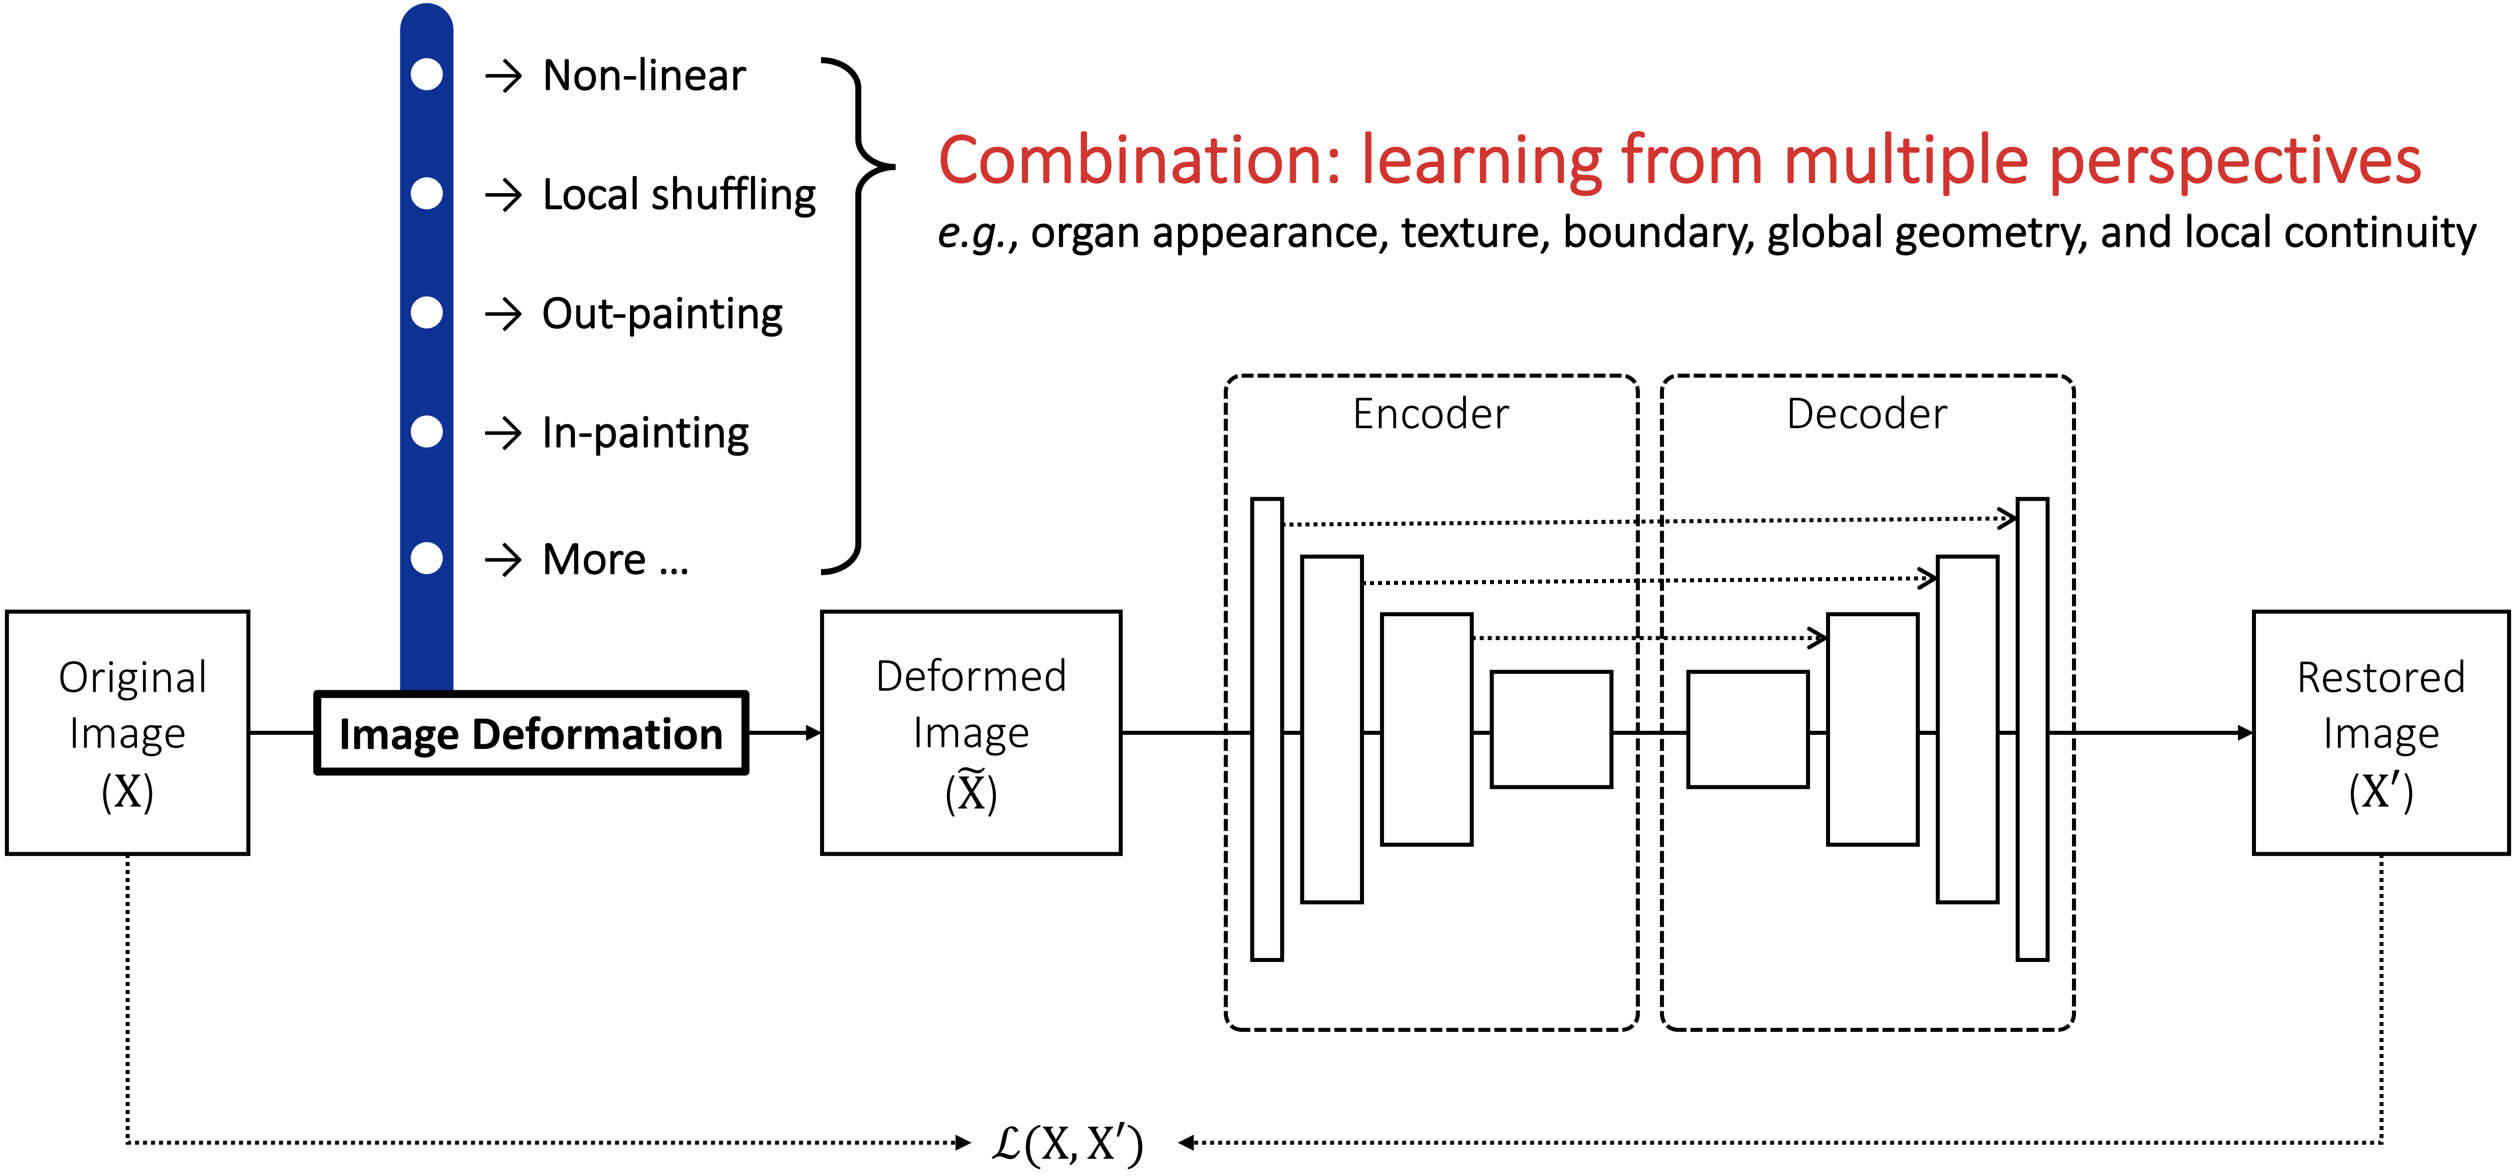
Non-linear
Local shuffling
Out-painting
In-painting
More …
Combination: learning from multiple perspectives
e.g., organ appearance, texture, boundary, global geometry, and local continuity
Original Image (X)
Image Deformation
Deformed Image ($\tilde{X}$)
Encoder
Decoder
Restored Image (X′)
$\mathcal{L}(X,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

Once pre-trained,

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

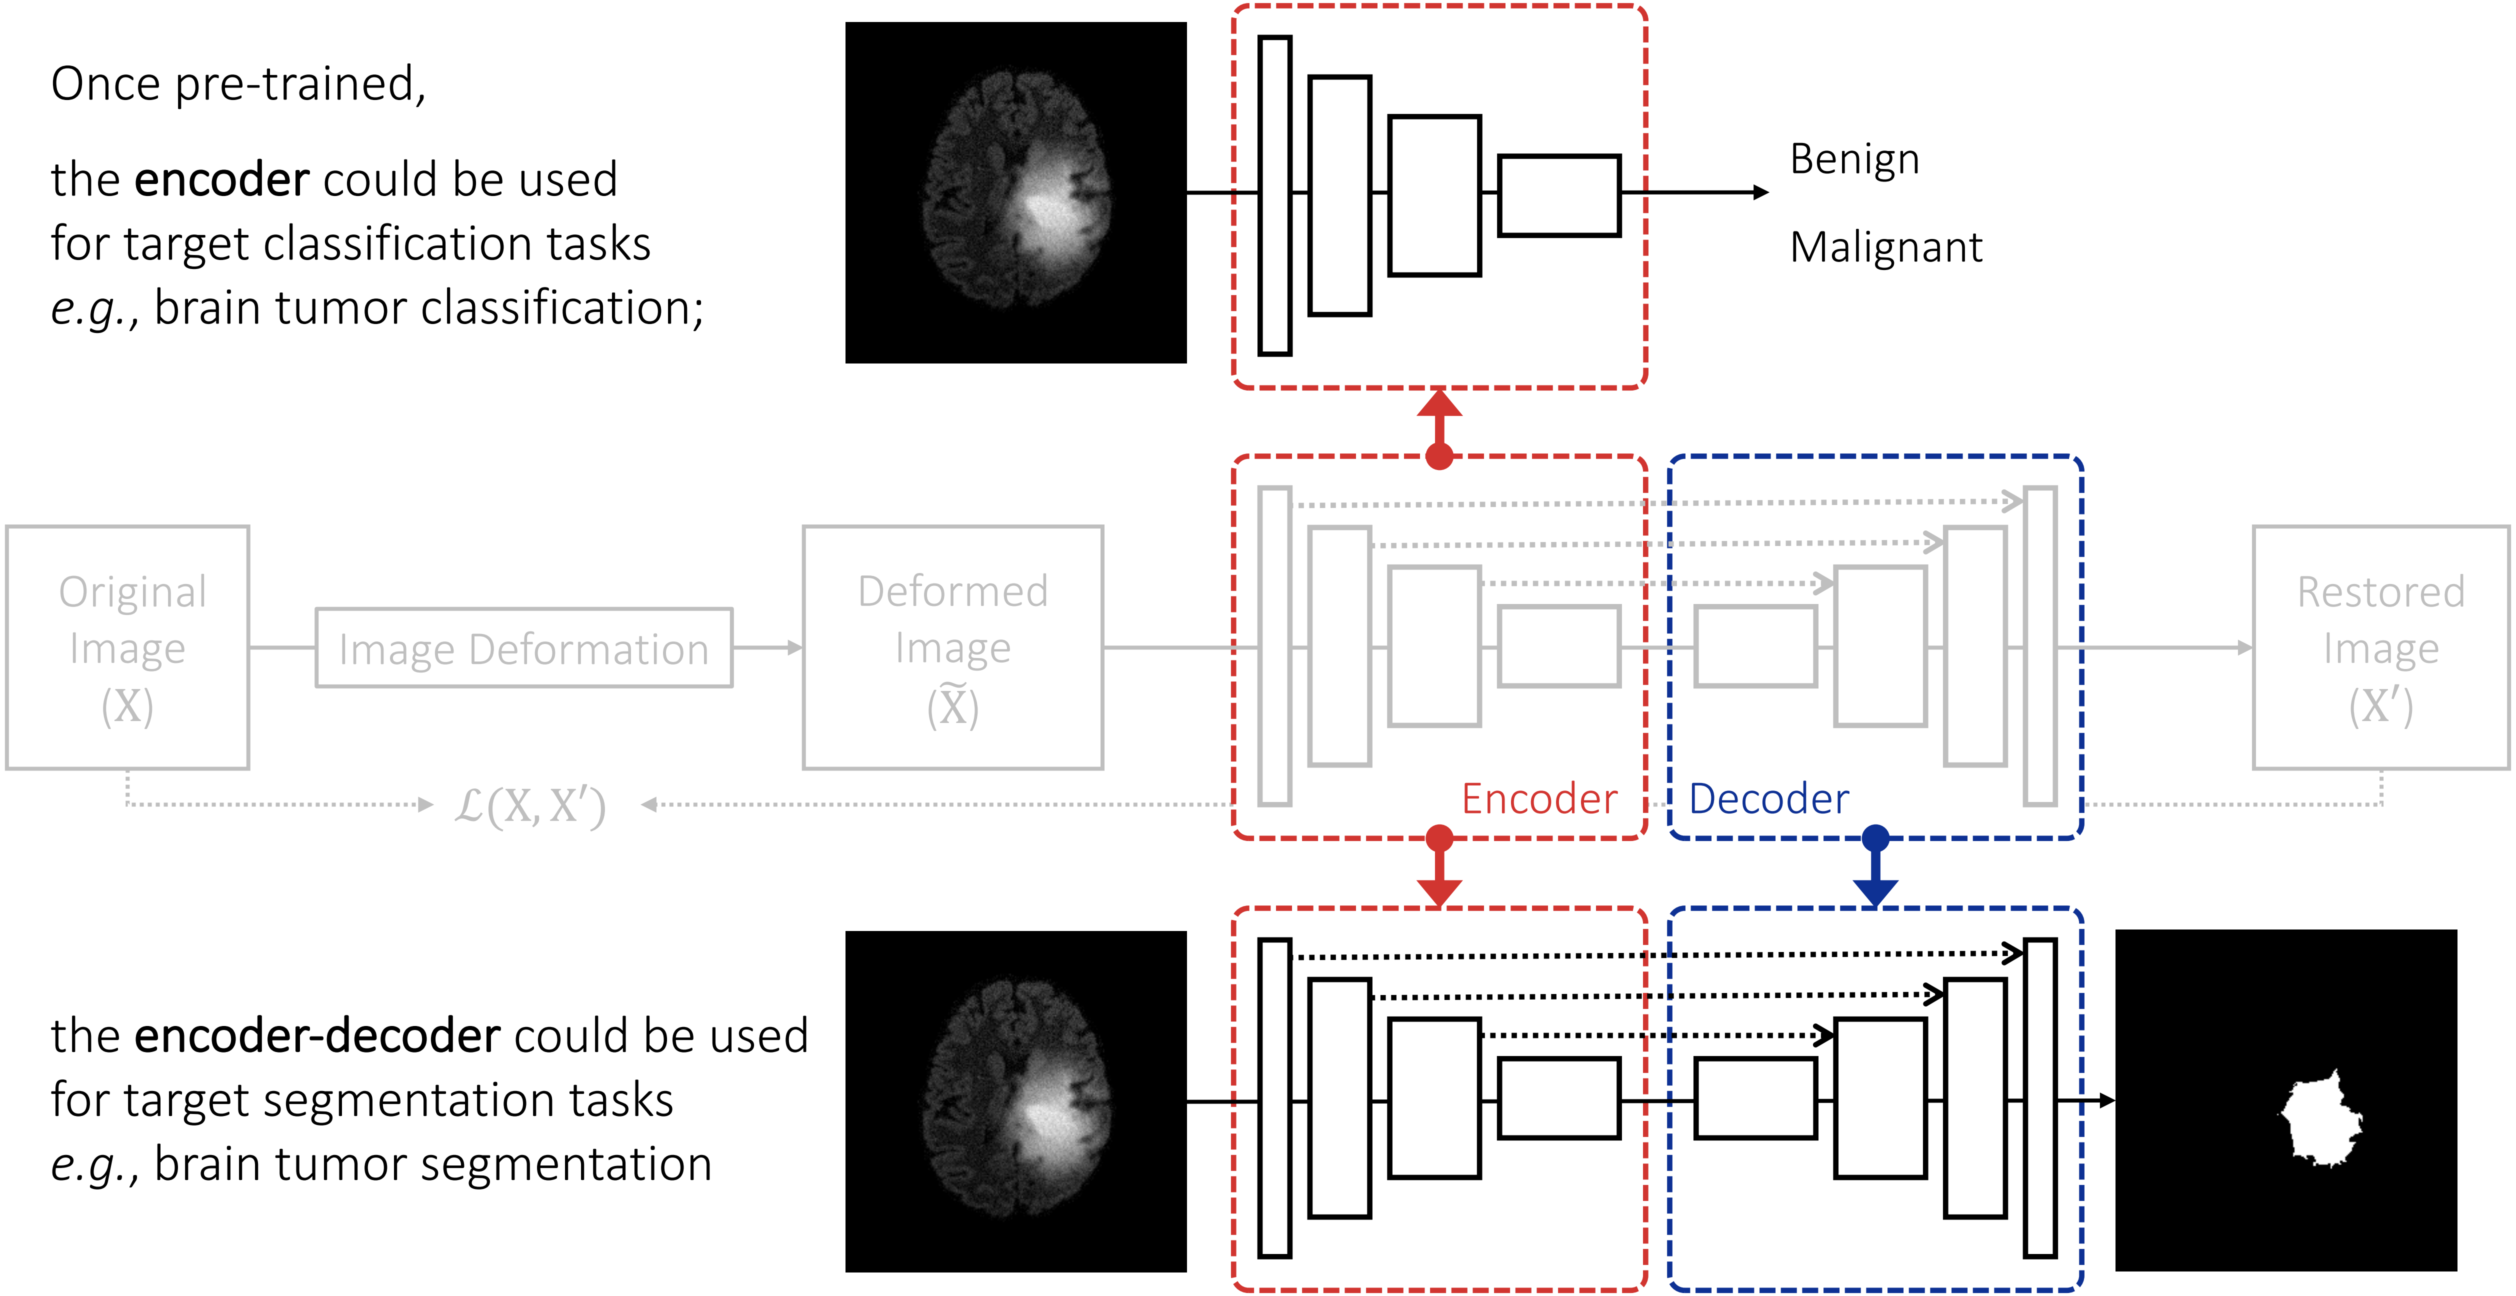

the **encoder-decoder** could be used
for target segmentation tasks
e.g., brain tumor segmentation

Four classification applications

1. Lung nodule false positive reduction (CT)
2. PE false positive reduction 3D CT (CT)
3. Eight pulmonary diseases classification (X-ray)
4. RoI/bulb/background classification (Ultrasound)

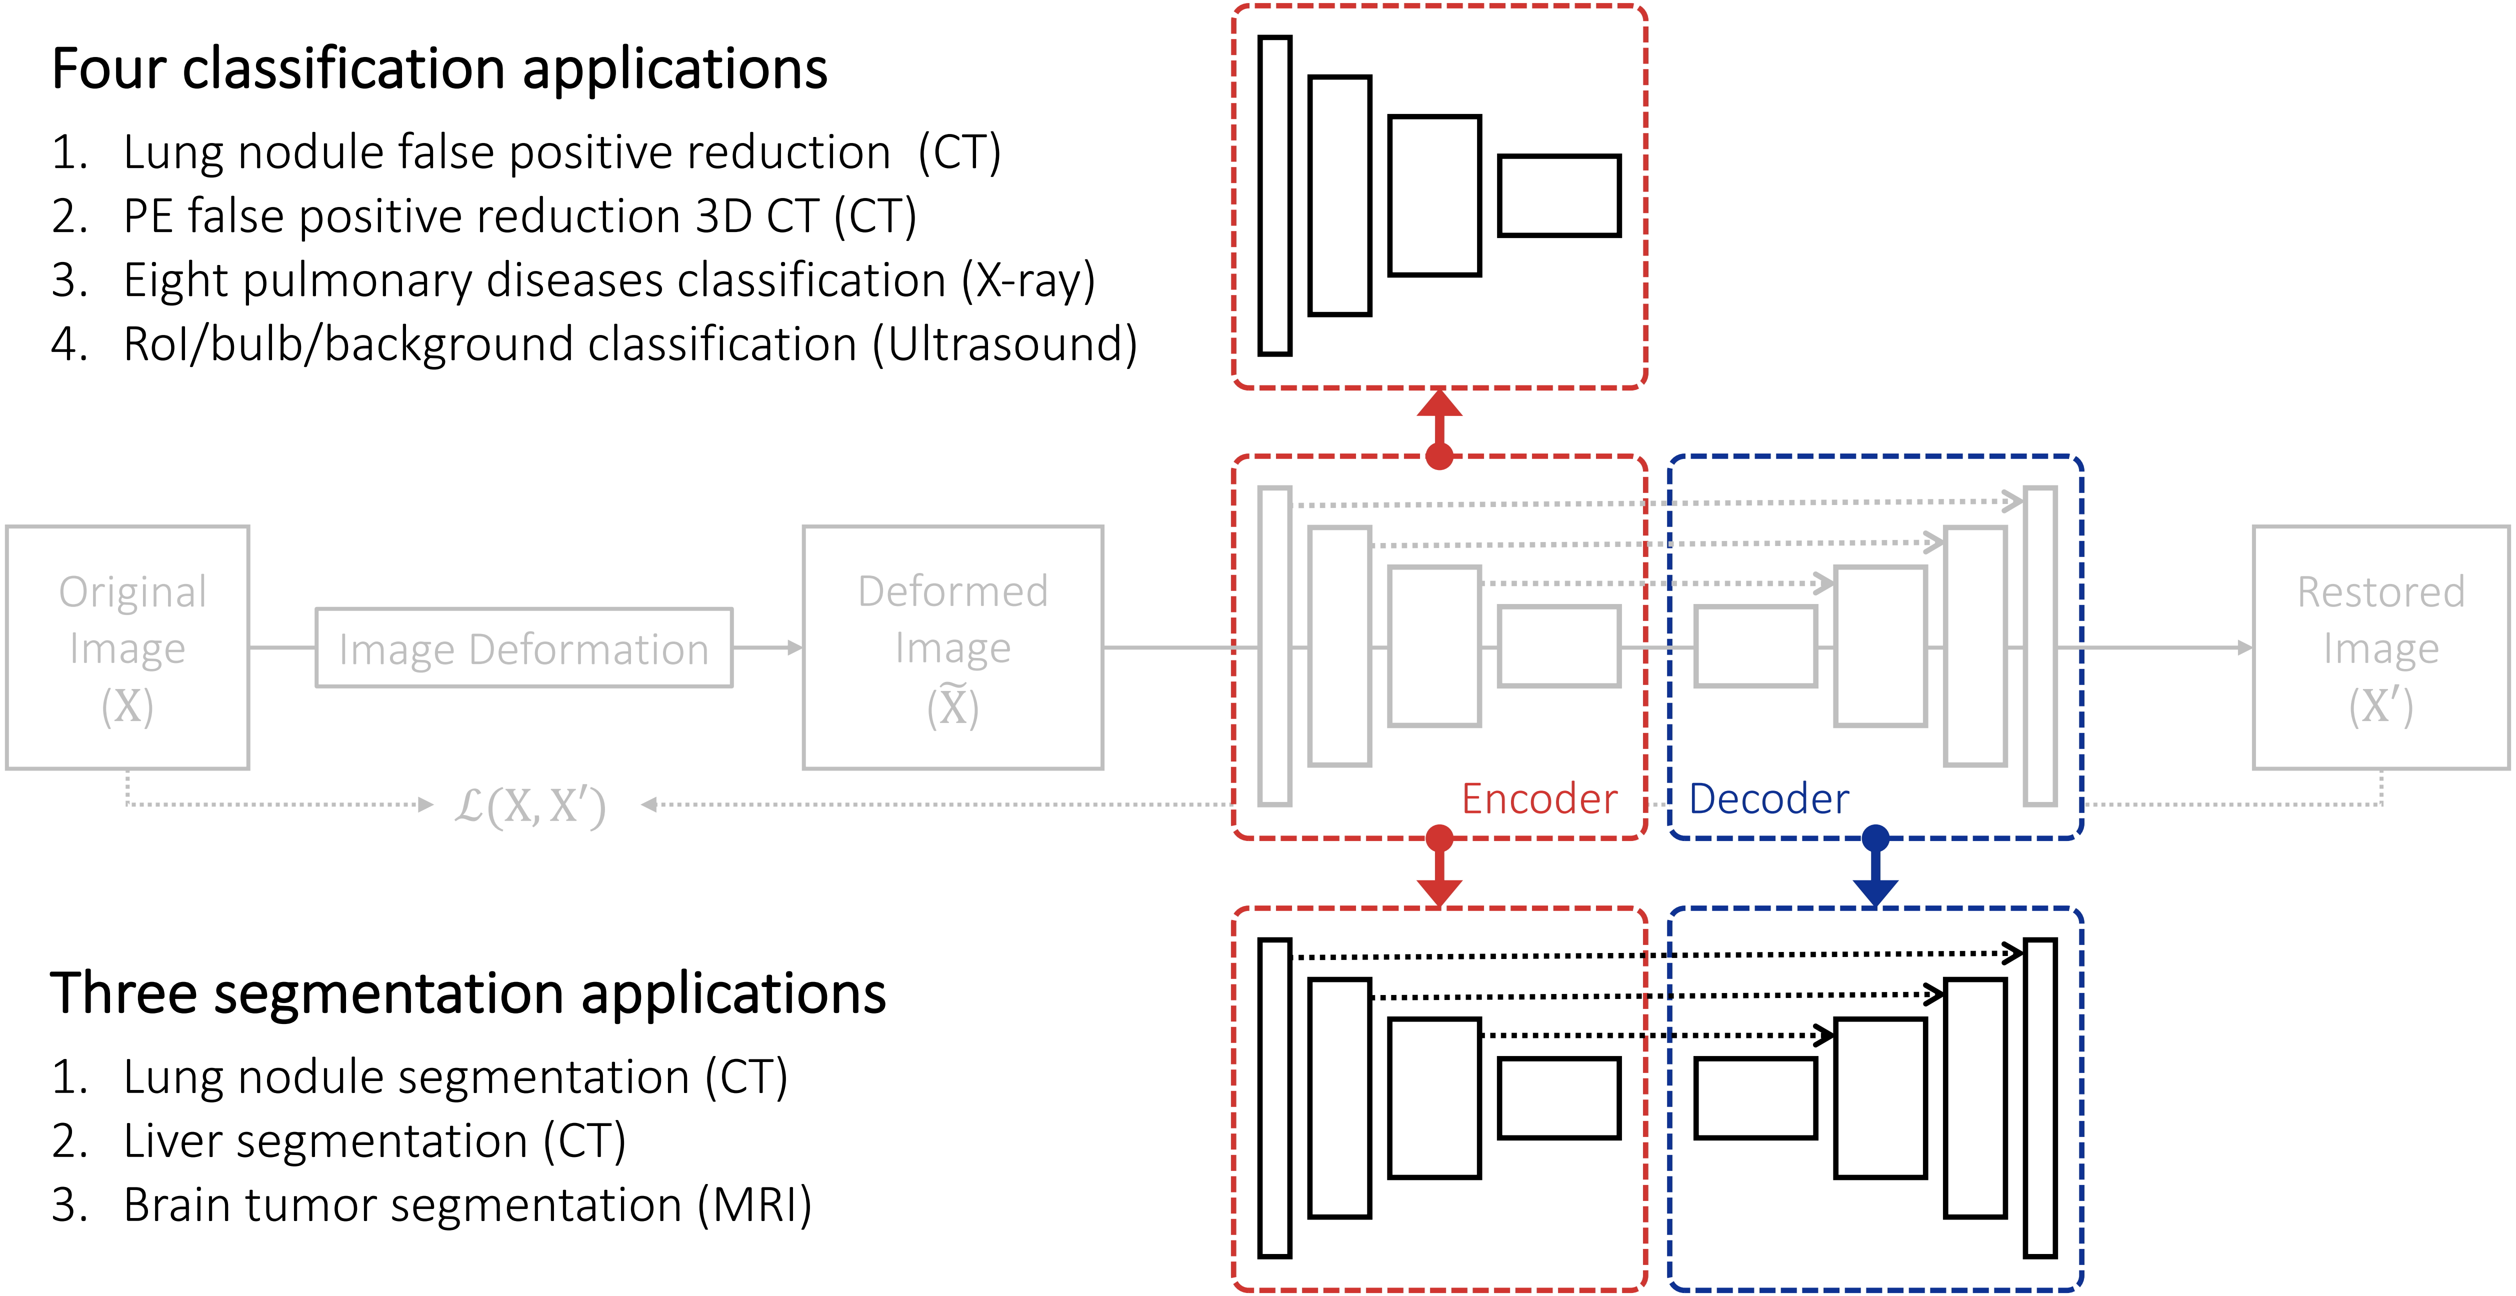

Three segmentation applications

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

Devils in 3D Models

: Learning from scratch *simply* in 3D may not necessarily yield performance better than fine-tuning from ImageNet in 2D

Scratch 2D
Scratch 3D
ImageNet
Genesis X-ray 2D
Genesis CT 2D
Genesis CT 3D

n.s.	No Significance	
*	$p < 0.05$	n=10
**	$p < 0.01$	n=10
***	$p < 0.001$	n=10
>***	$p < 0.0001$	n=10

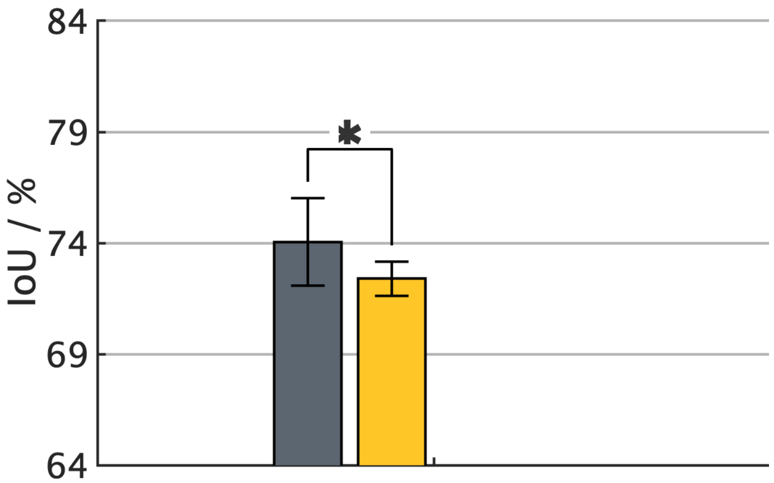

Lung nodule segmentation
3D CT [LIDC-IDR]

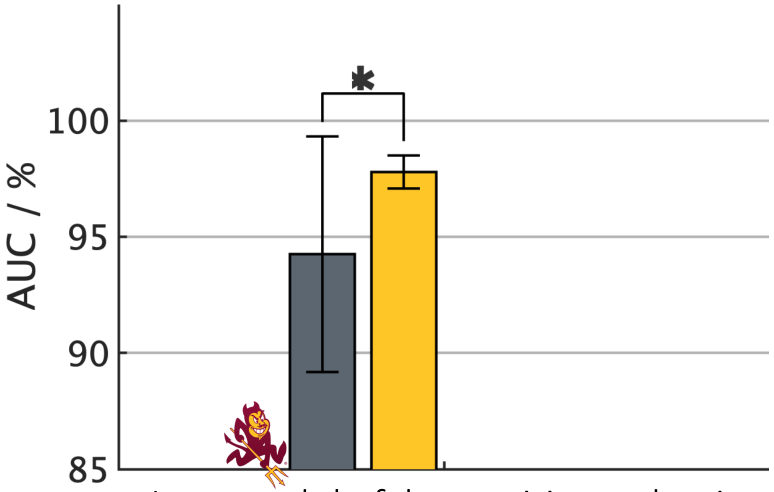

Lung nodule false positive reduction
3D CT [LUNA-2016]

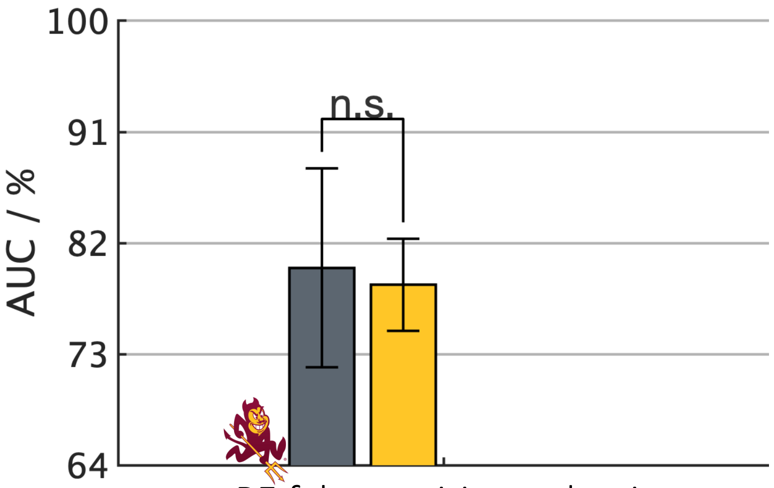

PE false positive reduction
3D CT [PE-CAD]

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

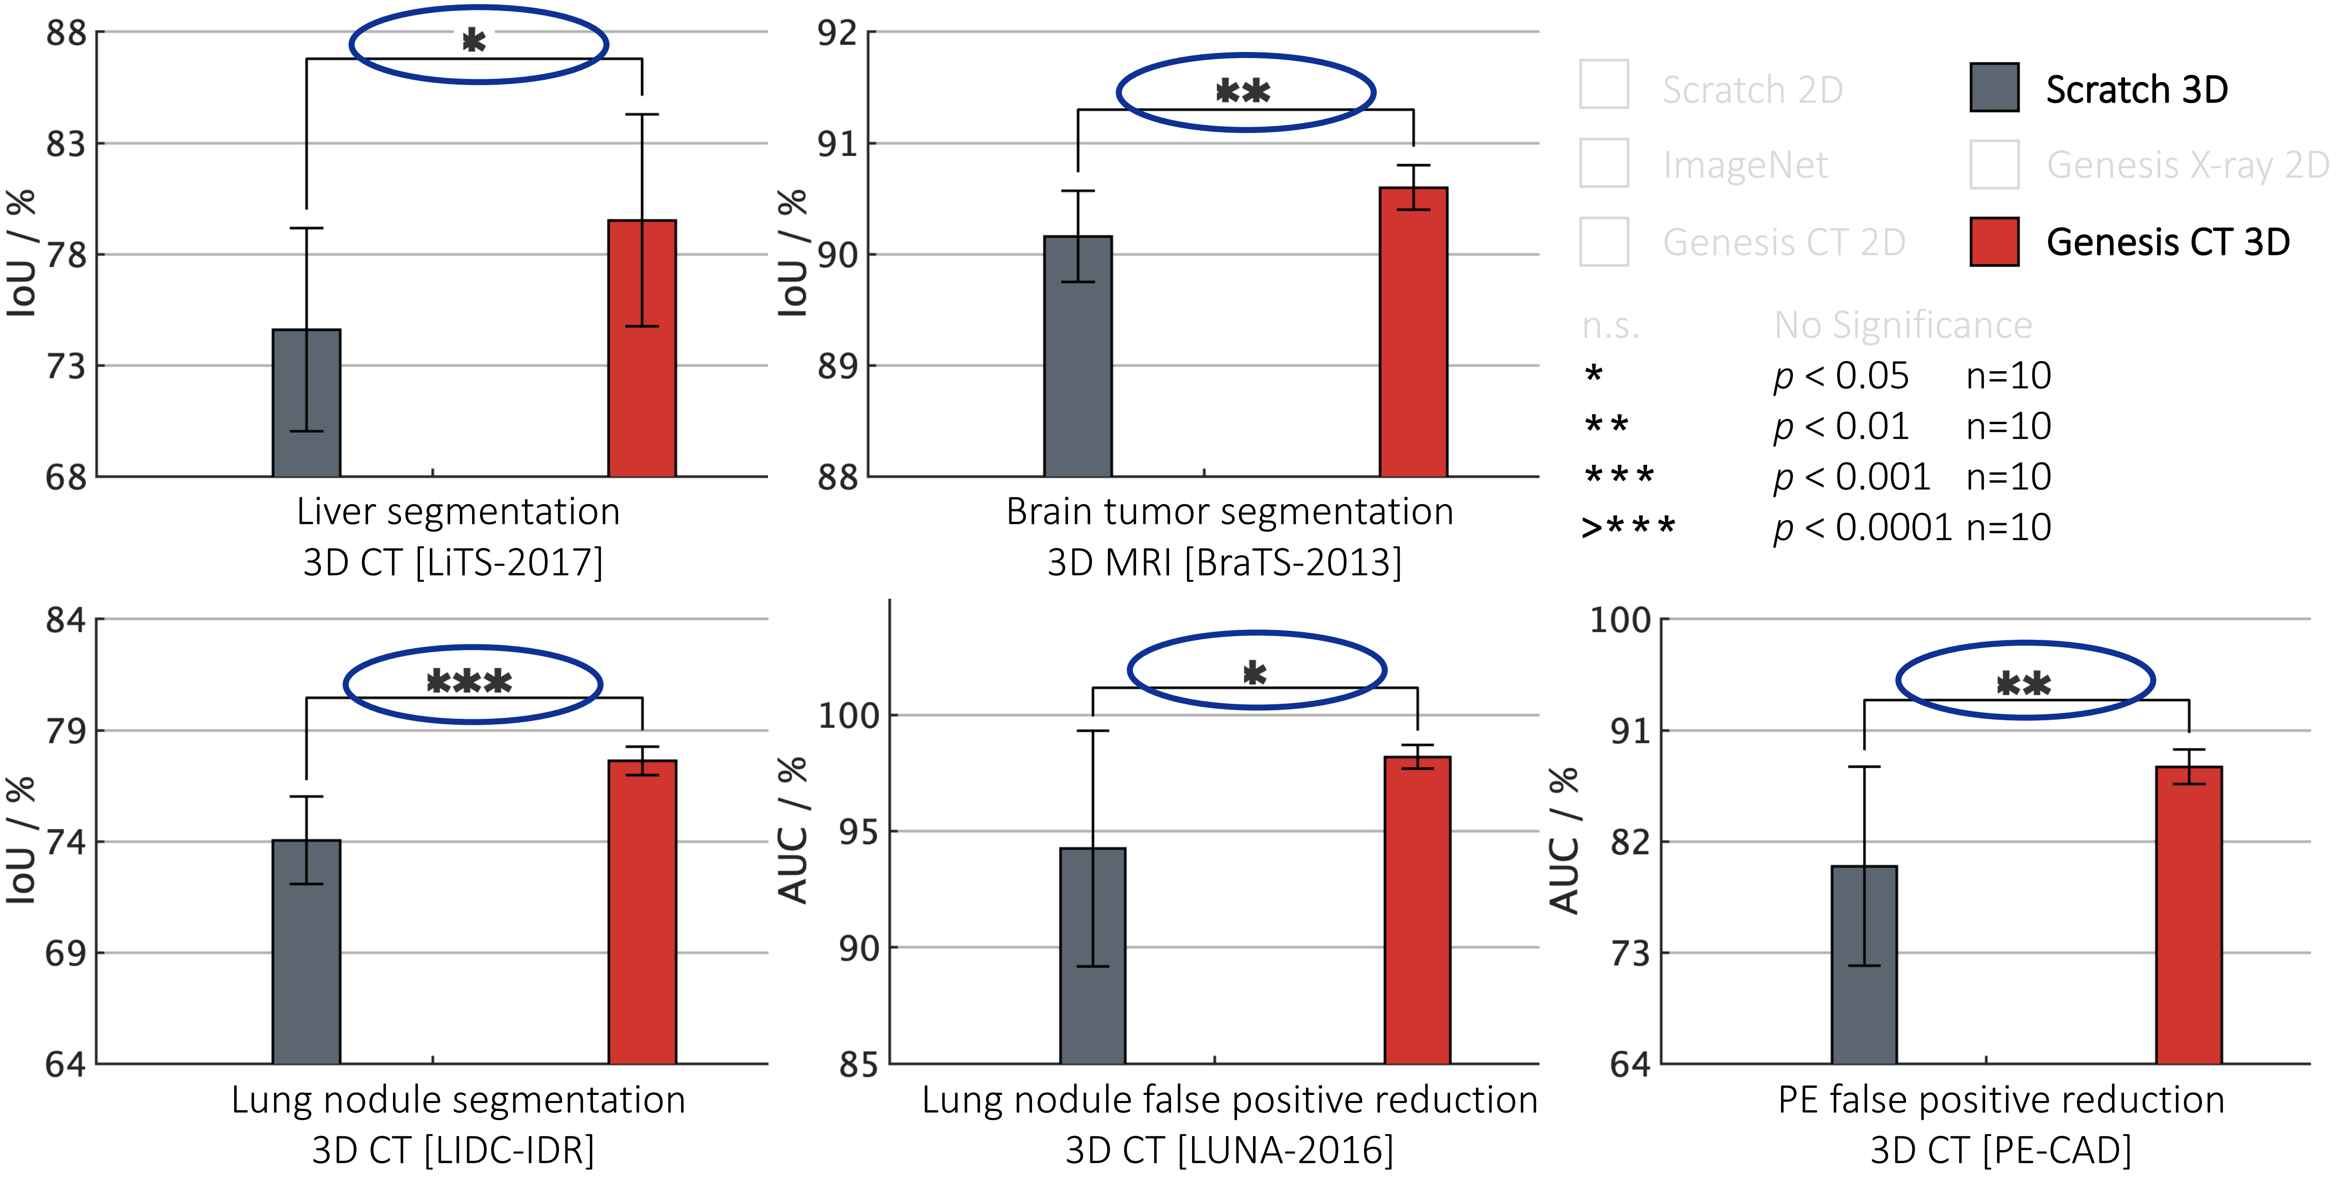

Result II: Models Genesis consistently outperform any 2D approaches

including

1. ImageNet (state-of-the-art)
2. Models Genesis 2D (degraded)
   - Genesis X-ray 2D: pre-trained on NIH X-ray dataset
   - Genesis CT 2D: pre-trained on LUNA-2016 dataset

Scratch 2D
Scratch 3D
ImageNet
Genesis X-ray 2D
Genesis CT 2D
Genesis CT 3D

n.s.	No Significance	
*	$p < 0.05$	n=10
**	$p < 0.01$	n=10
***	$p < 0.001$	n=10
>***	$p < 0.0001$	n=10

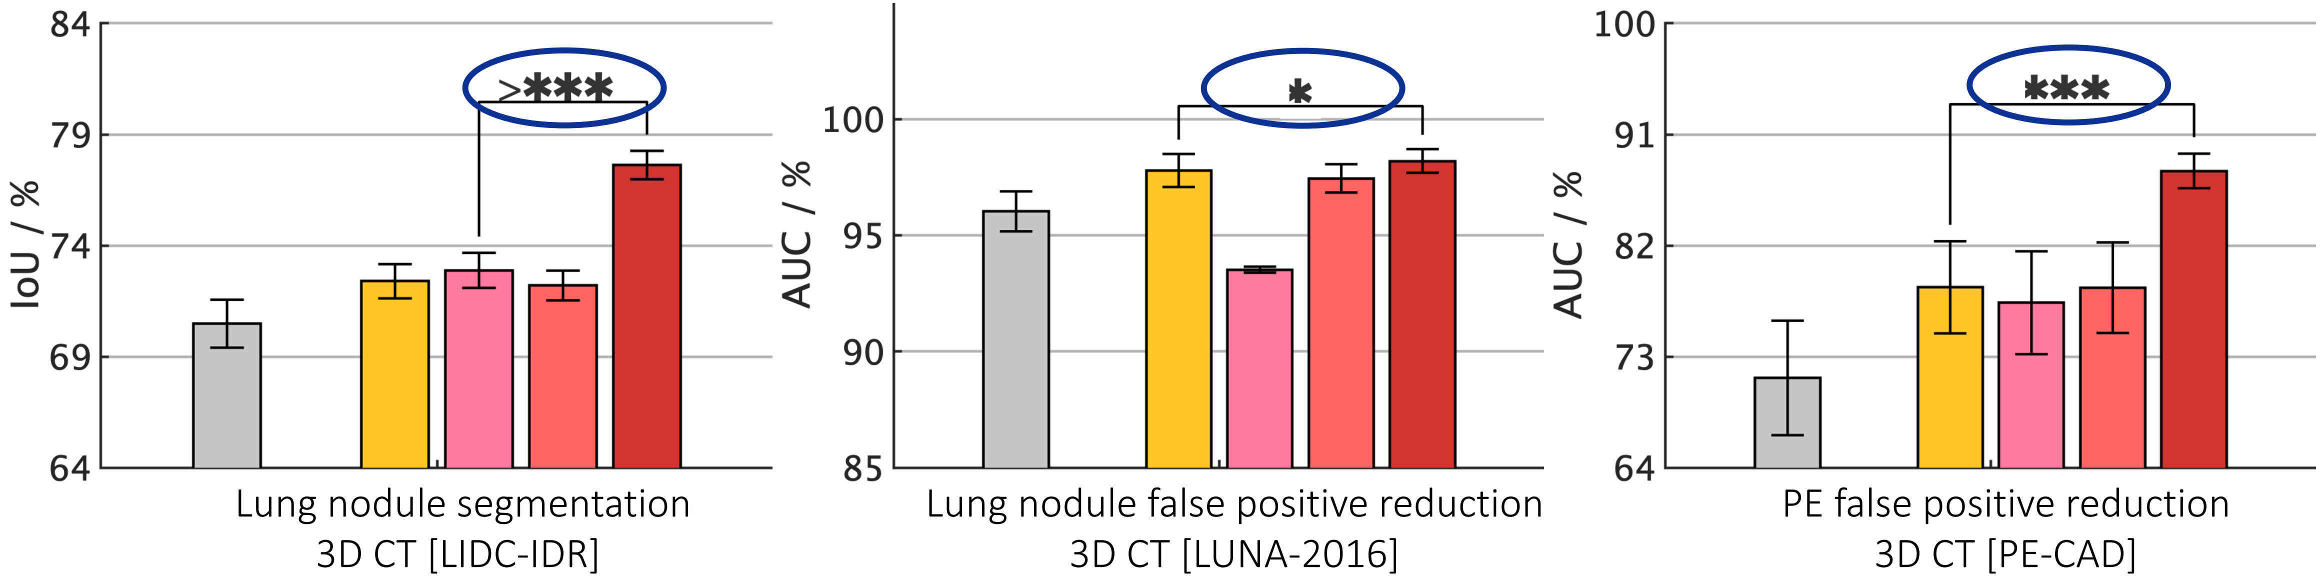

Result III: Models Genesis 2D (self-supervised) ≈ ImageNet (supervised)

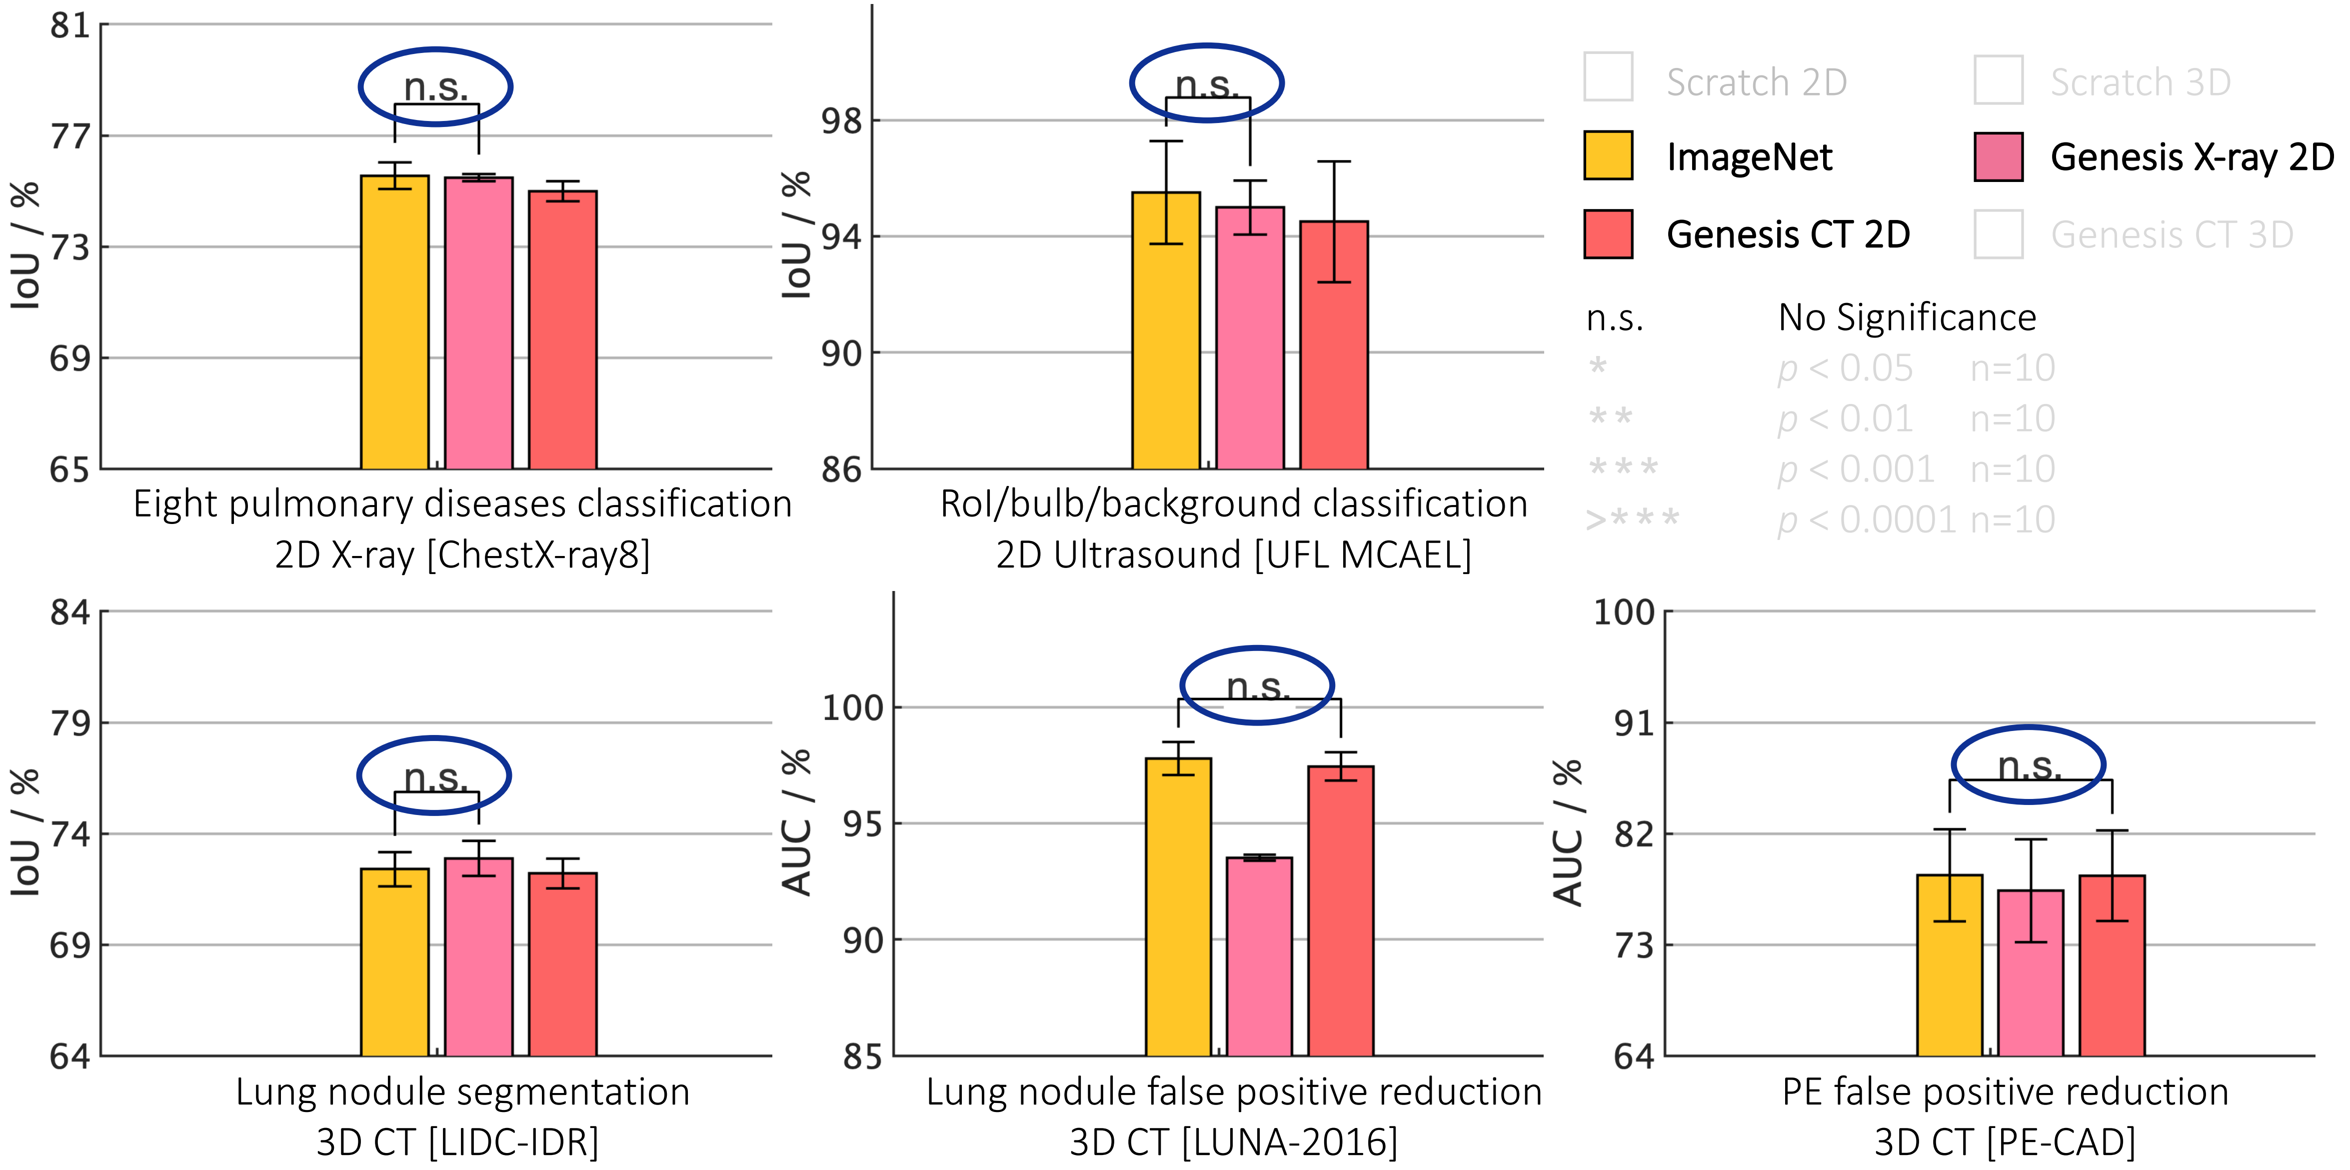

Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis

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

1. Models Genesis outperform 3D models trained from scratch
2. Models Genesis consistently outperform any 2D approaches
3. Models Genesis (2D) offer performances equivalent to supervised pre-trained models

Genesis Chest CT

Genesis X-ray

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

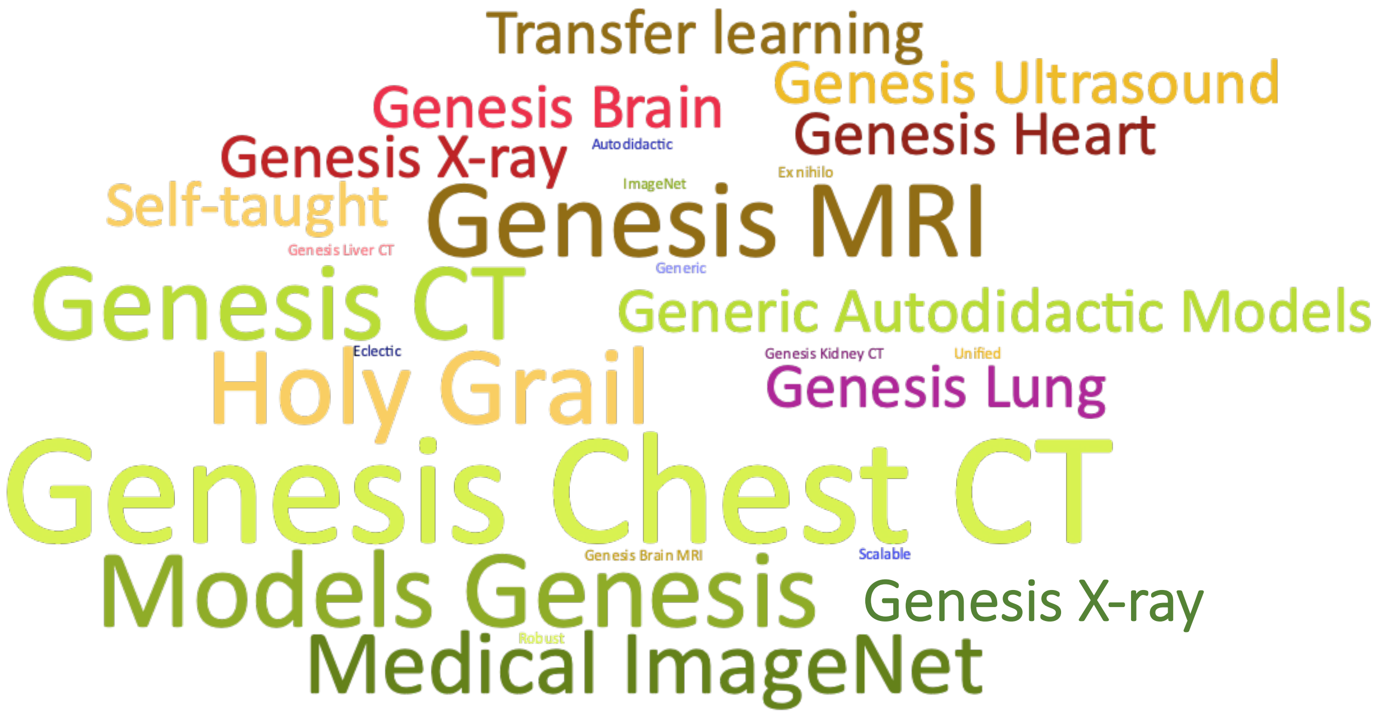

Models Genesis: Generic Autodidactic Models for 3D Medical Image Analysis

Poster session

Tuesday, October 15
13:00 – 14:00

Poster T-5-B-013

Try it for yourself

Code, data, and models are available online

github.com/MrGiovanni/ModelsGenesis