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。”

——沈定刚教授

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

Important Links

Code and pre-trained weights

Full paper and appendix

Oral presentation

Download poster in PDF

中文博客

Download slides in PDF

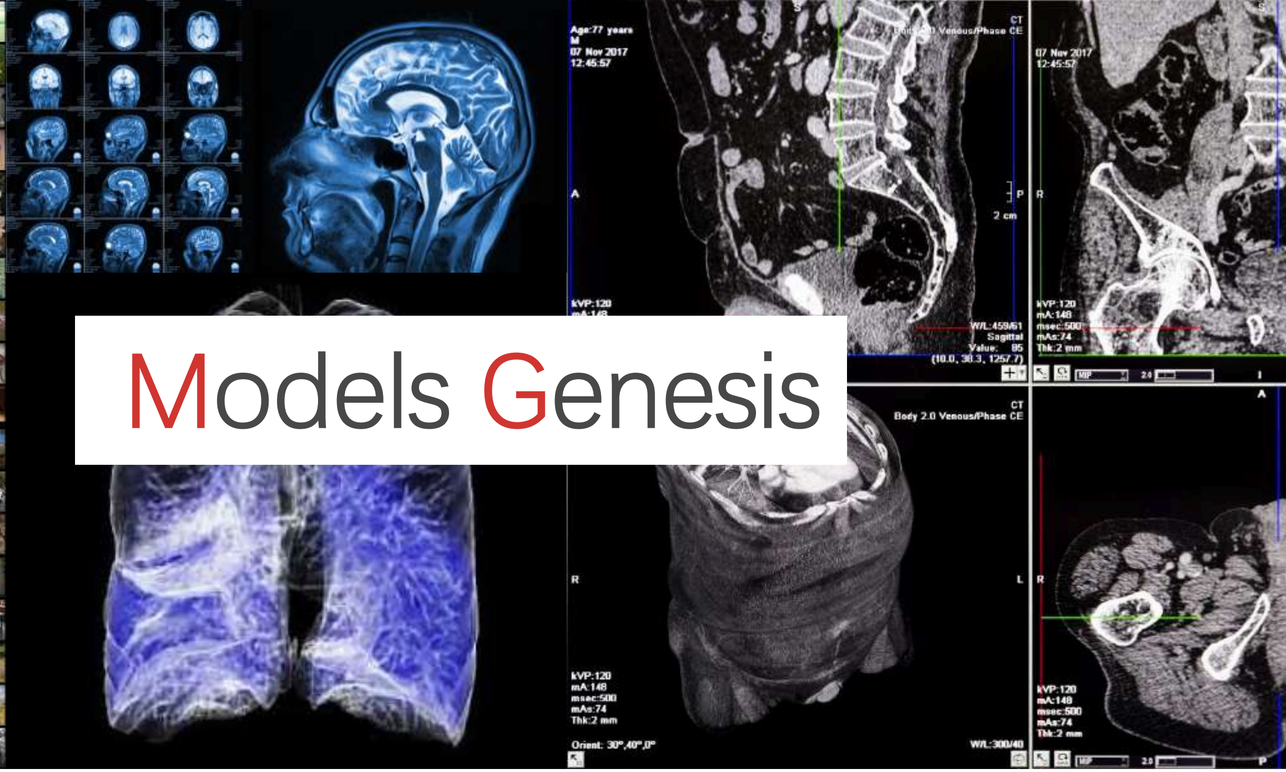

Natural images	Medical images
Formed in 2D	Formed in 3D
>14,000,000 annotation	Zero annotation

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

Self-supervised learning

AI研习社

To learn representation directly from data itself

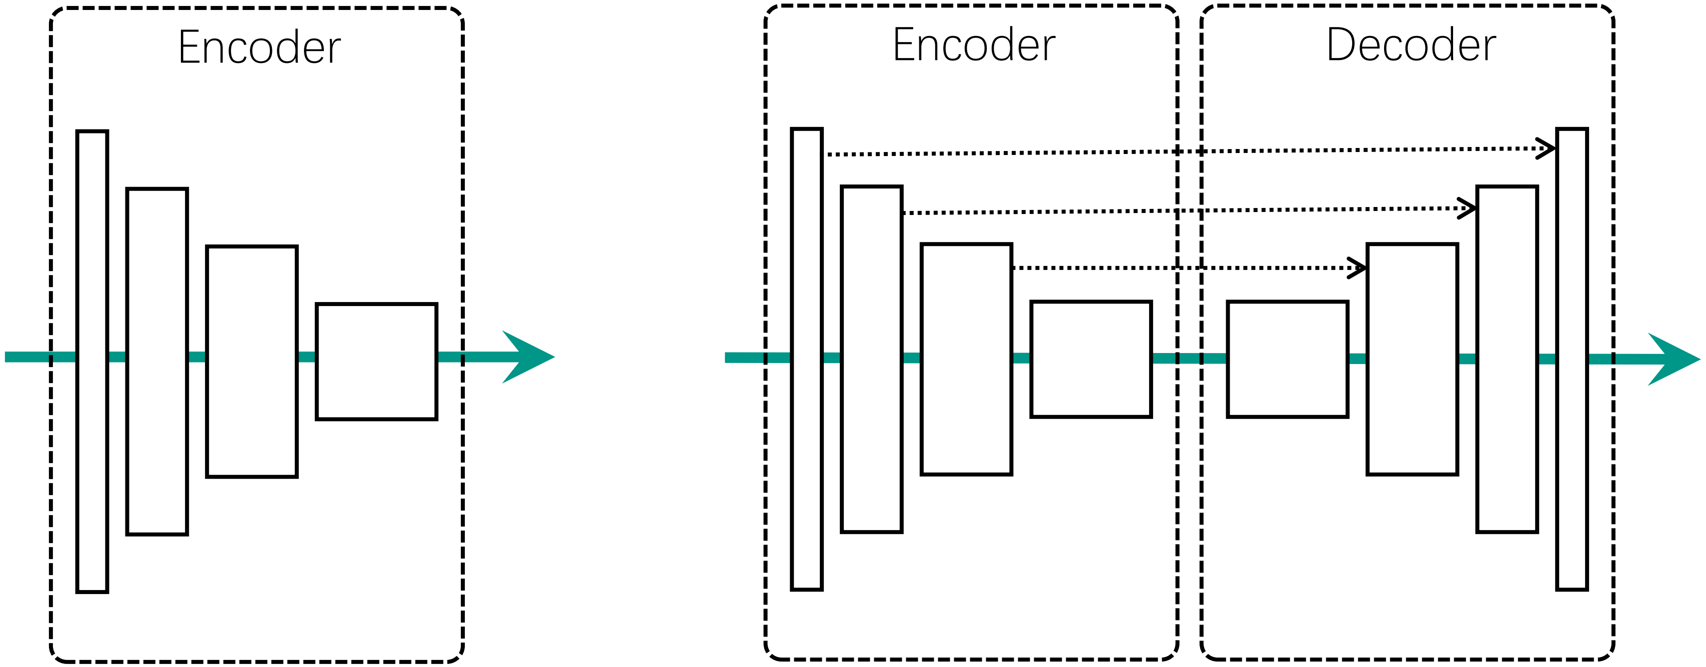

Zisserman et al. "Self-Supervised Learning." https://project.inria.fr/paiss/files/2018/07/zisserman-self-supervised.pdf
Ren et al. "Awesome Self-Supervised Learning." https://github.com/jason718/awesome-self-supervised-learning
Jing, et al. "Self-supervised Visual Feature Learning with Deep Neural Networks: A Survey." arXiv preprint arXiv:1902.06162 (2019). https://arxiv.org/pdf/1902.06162.pdf

Self-supervised learning

To learn representation directly from data itself

Example #1: patch relative positions

What is wrong when applying to medical imaging?

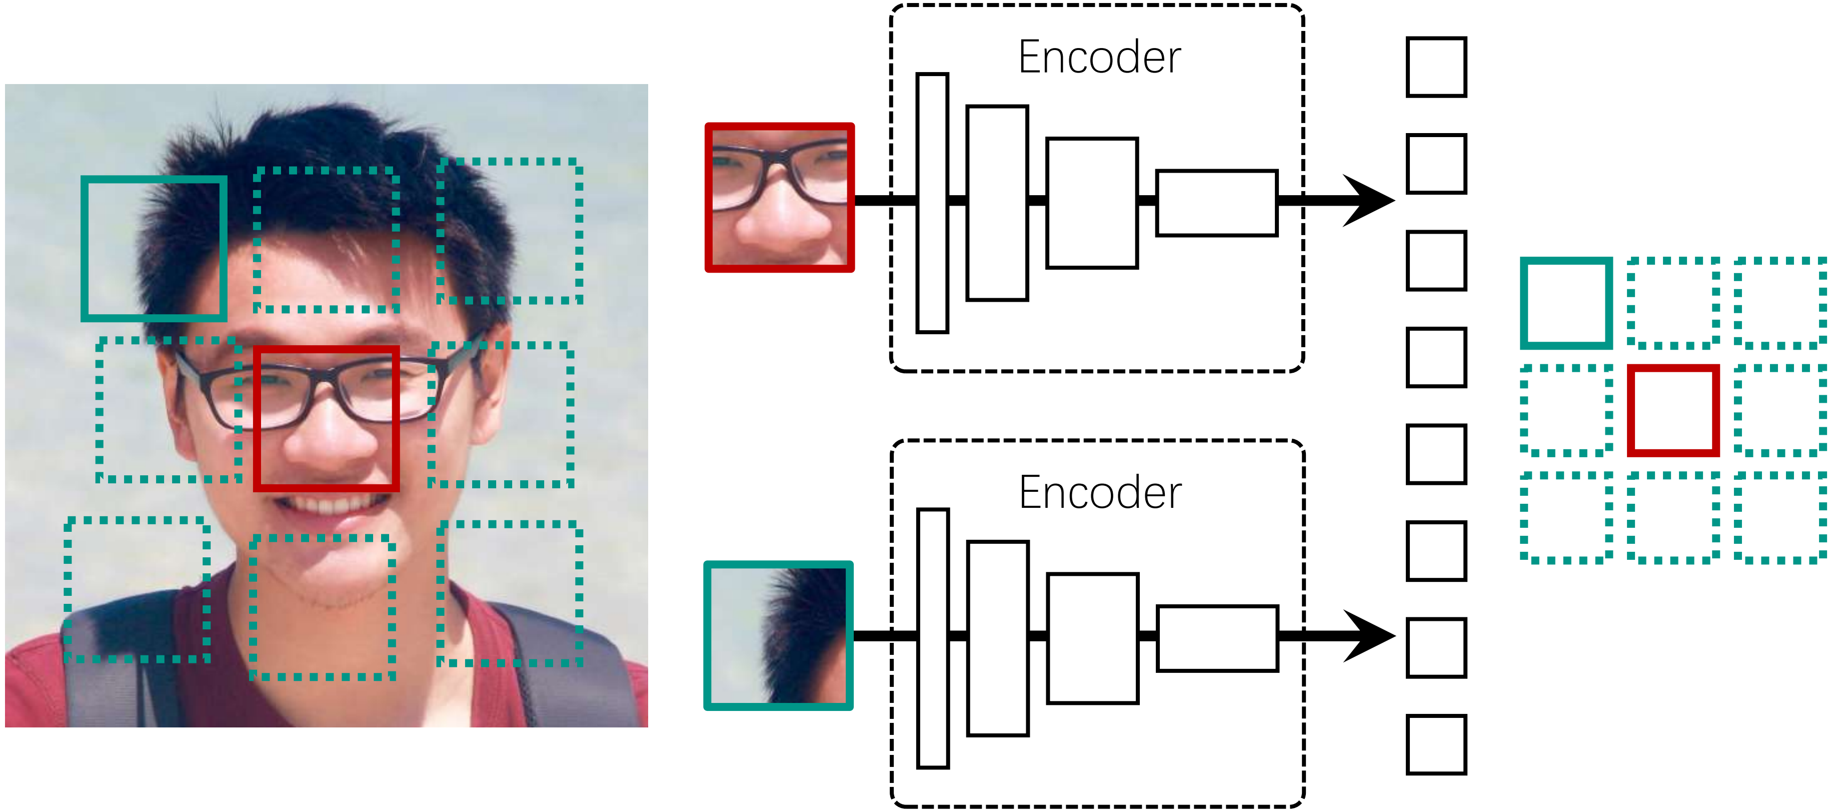

Doersch, et al. "Unsupervised visual representation learning by context prediction." In Proceedings of the IEEE International Conference on Computer Vision (2015).

Self-supervised learning

To learn representation directly from data itself

Example #1: patch relative positions

What is wrong when applying to medical imaging?

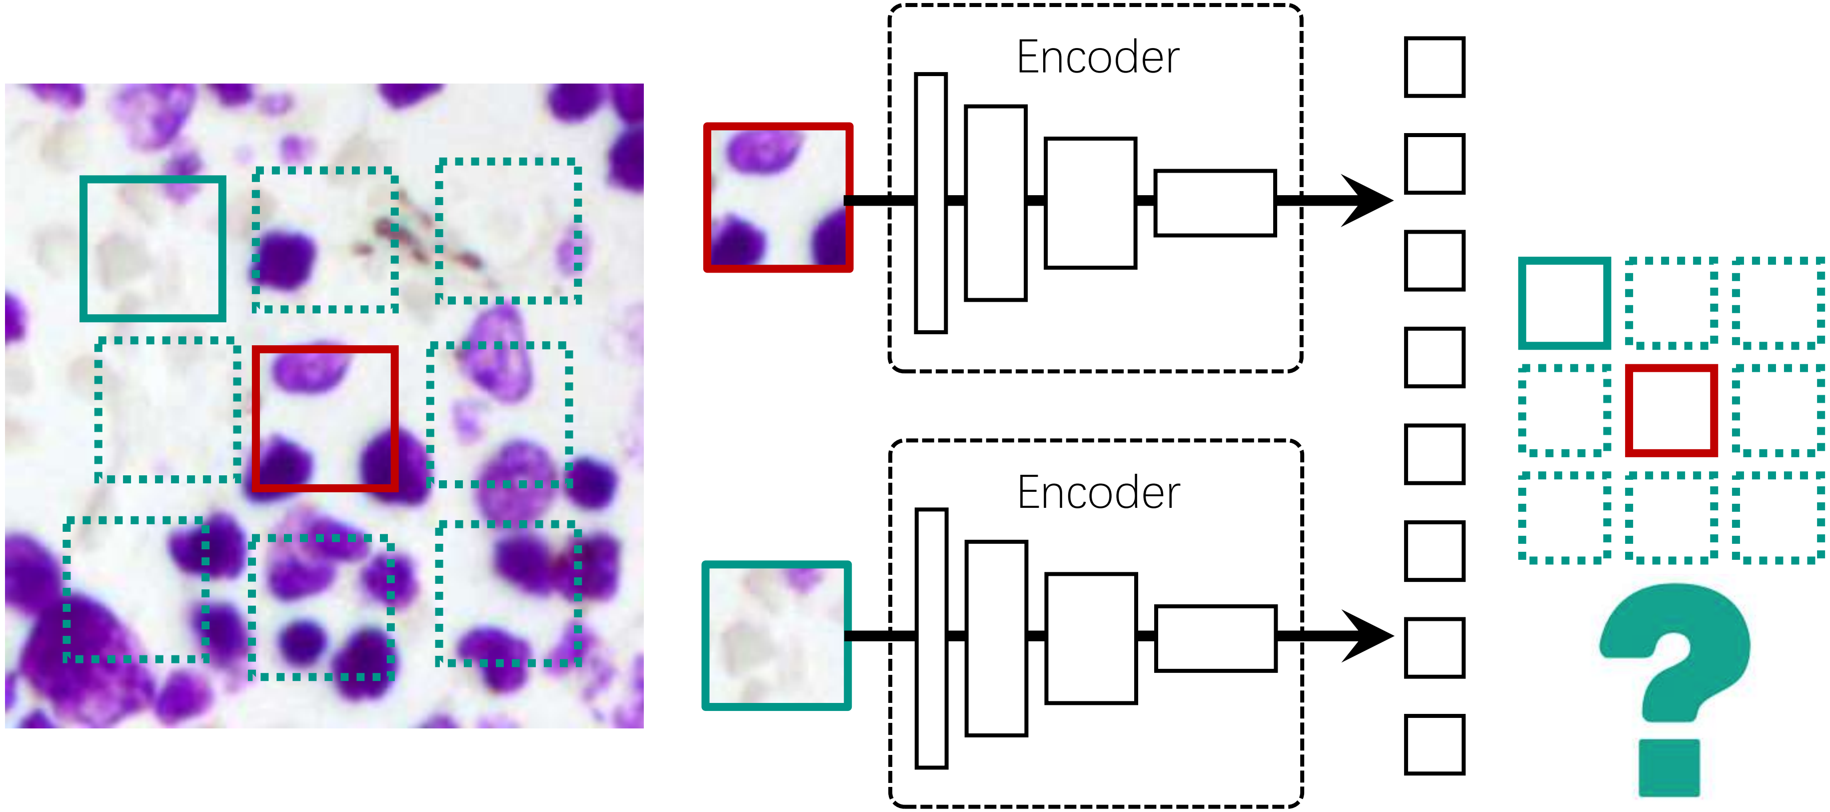

Doersch, et al. "Unsupervised visual representation learning by context prediction." In Proceedings of the IEEE International Conference on Computer Vision (2015).

Self-supervised learning

To learn representation directly from data itself

Example #2: image rotation

What is wrong when applying to medical imaging?

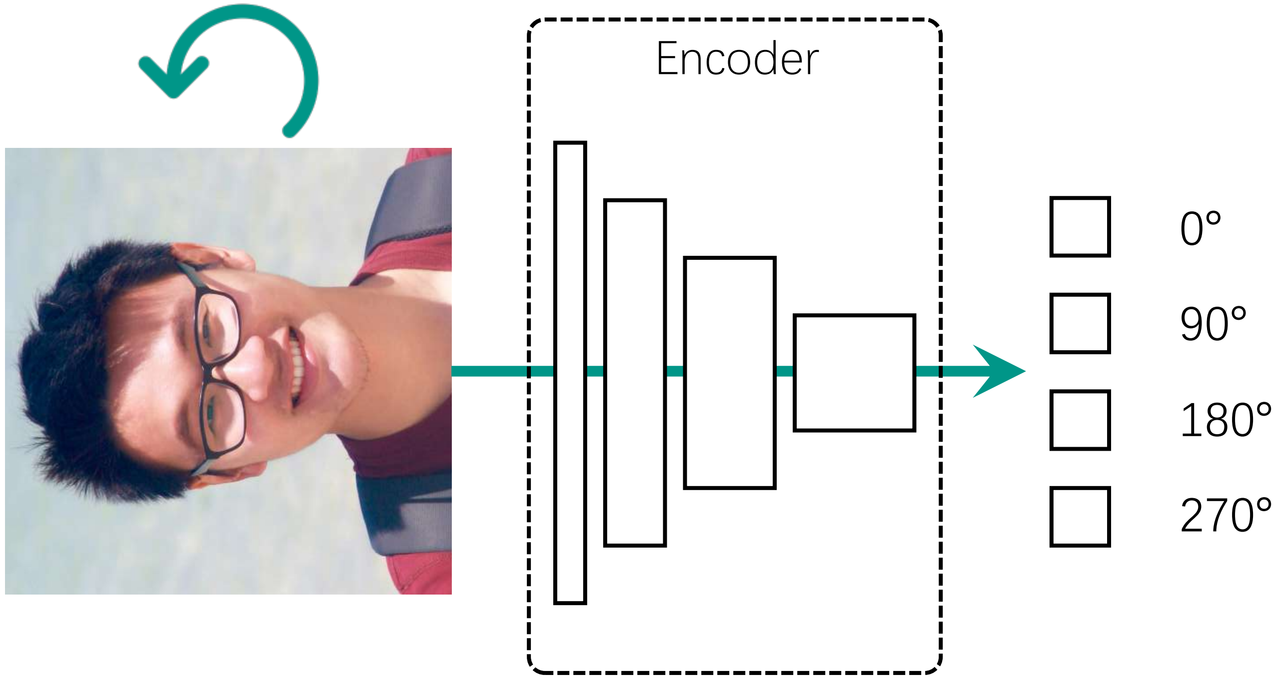

Gidaris, et al. "Unsupervised representation learning by predicting image rotations." In International Conference on Learning Representations (2018).

Self-supervised learning

To learn representation directly from data itself

Example #2: image rotation

What is wrong when applying to medical imaging?

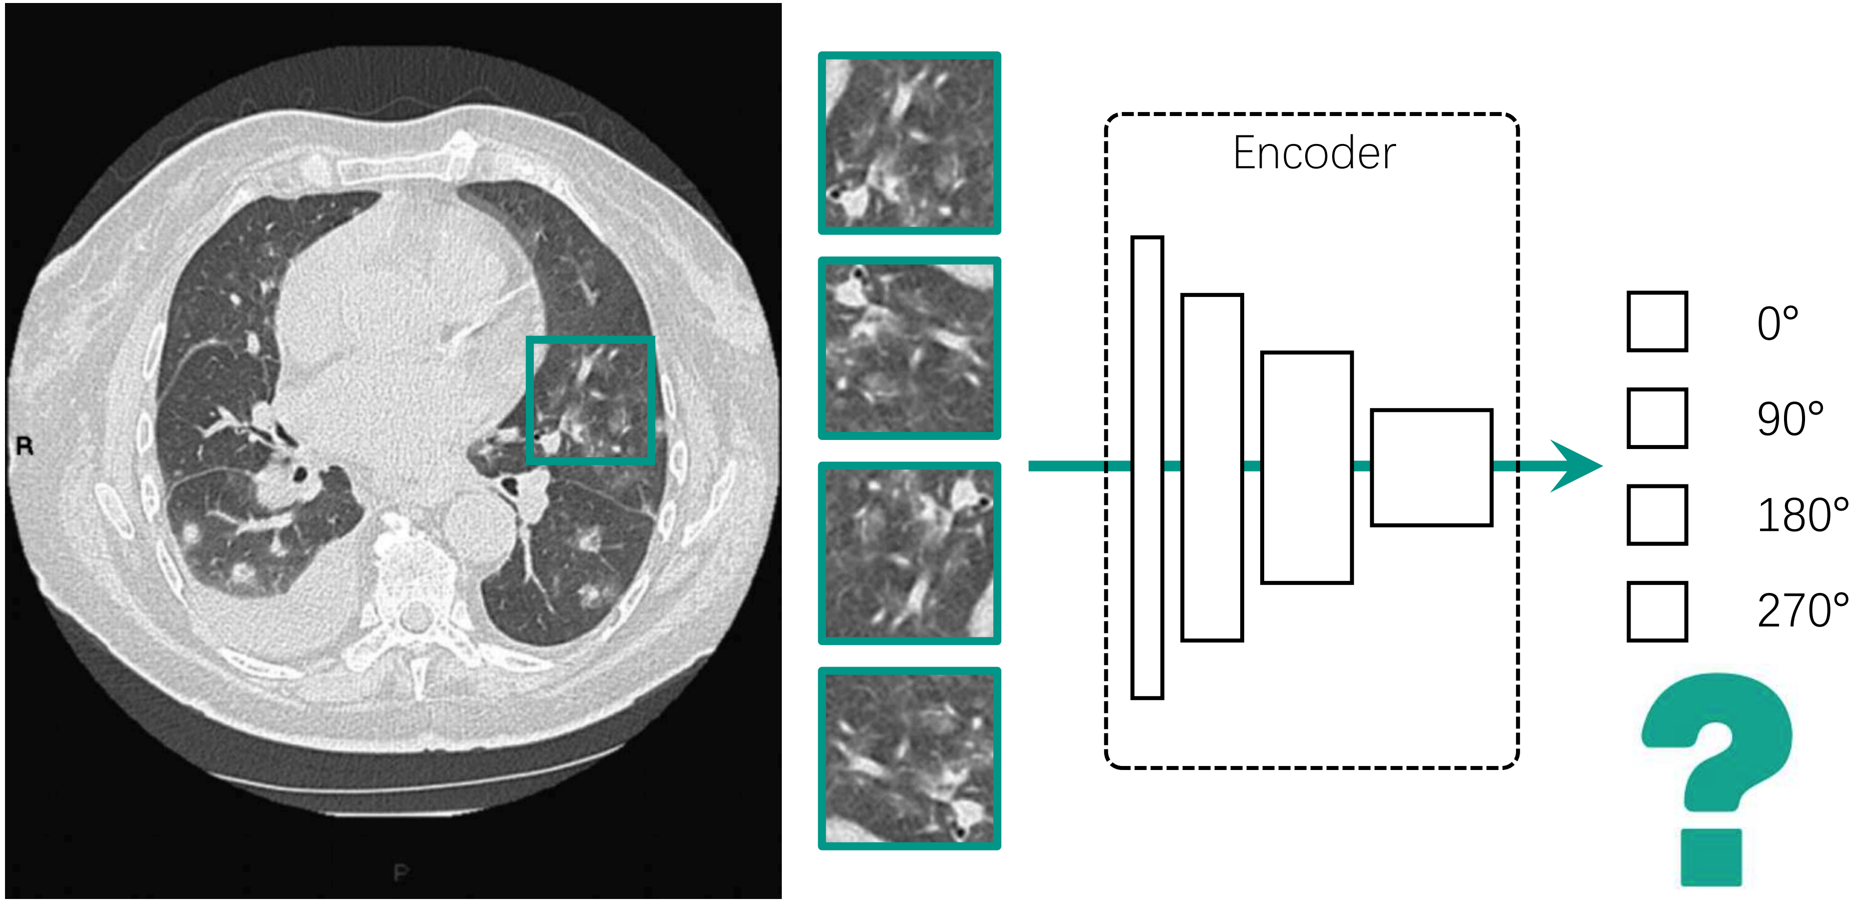

Gidaris, et al. "Unsupervised representation learning by predicting image rotations." In International Conference on Learning Representations (2018).

Self-supervised learning

To learn representation directly from data itself

Example #3: image colorization

What is wrong when applying to medical imaging?

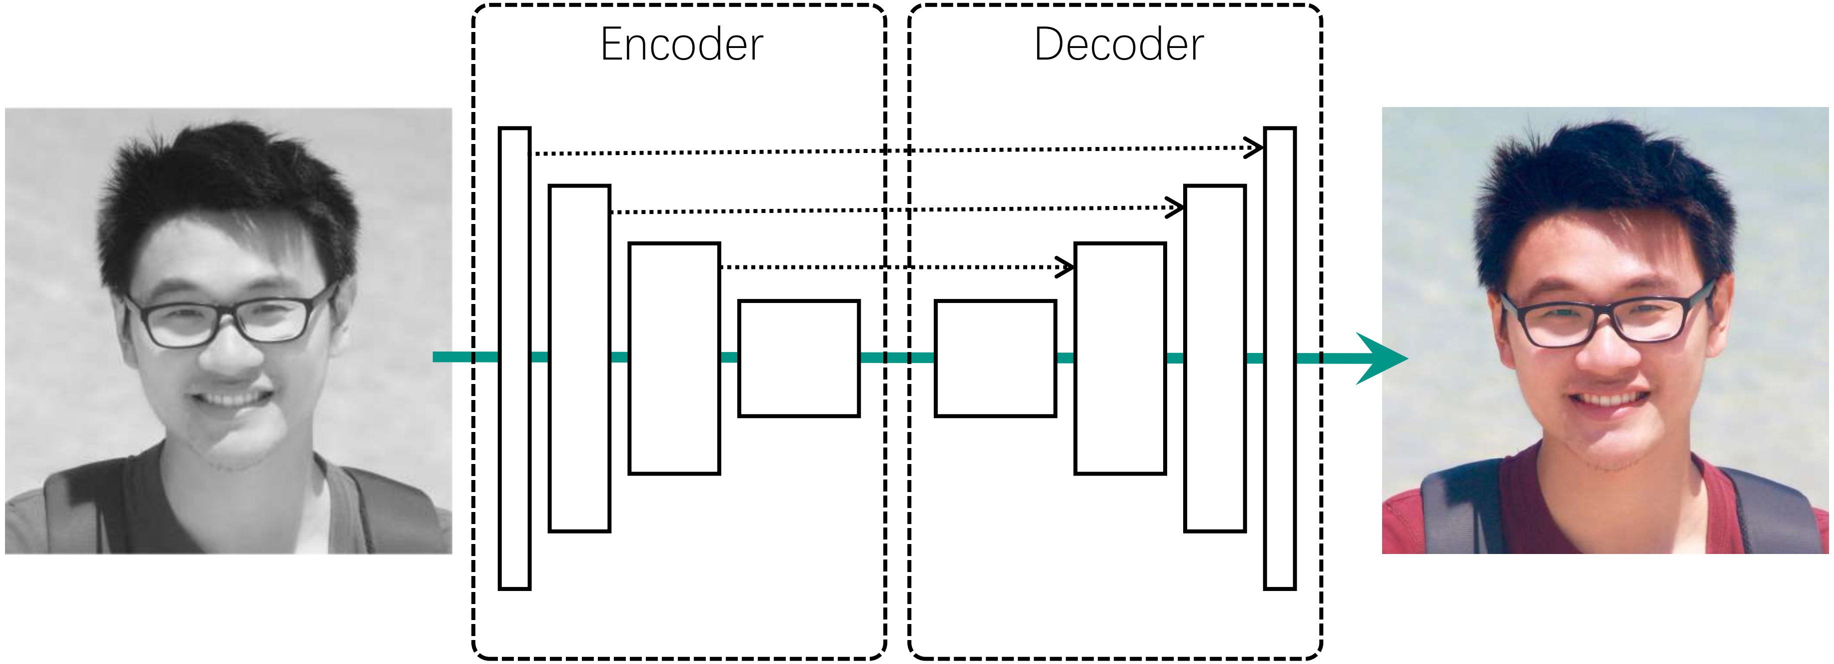

Zhang, et al. "Colorful image colorization." In European Conference on Computer Vision (2016).

AI研习社

Self-supervised learning

To learn representation directly from data itself

Example #3: image colorization

What is wrong when applying to medical imaging?

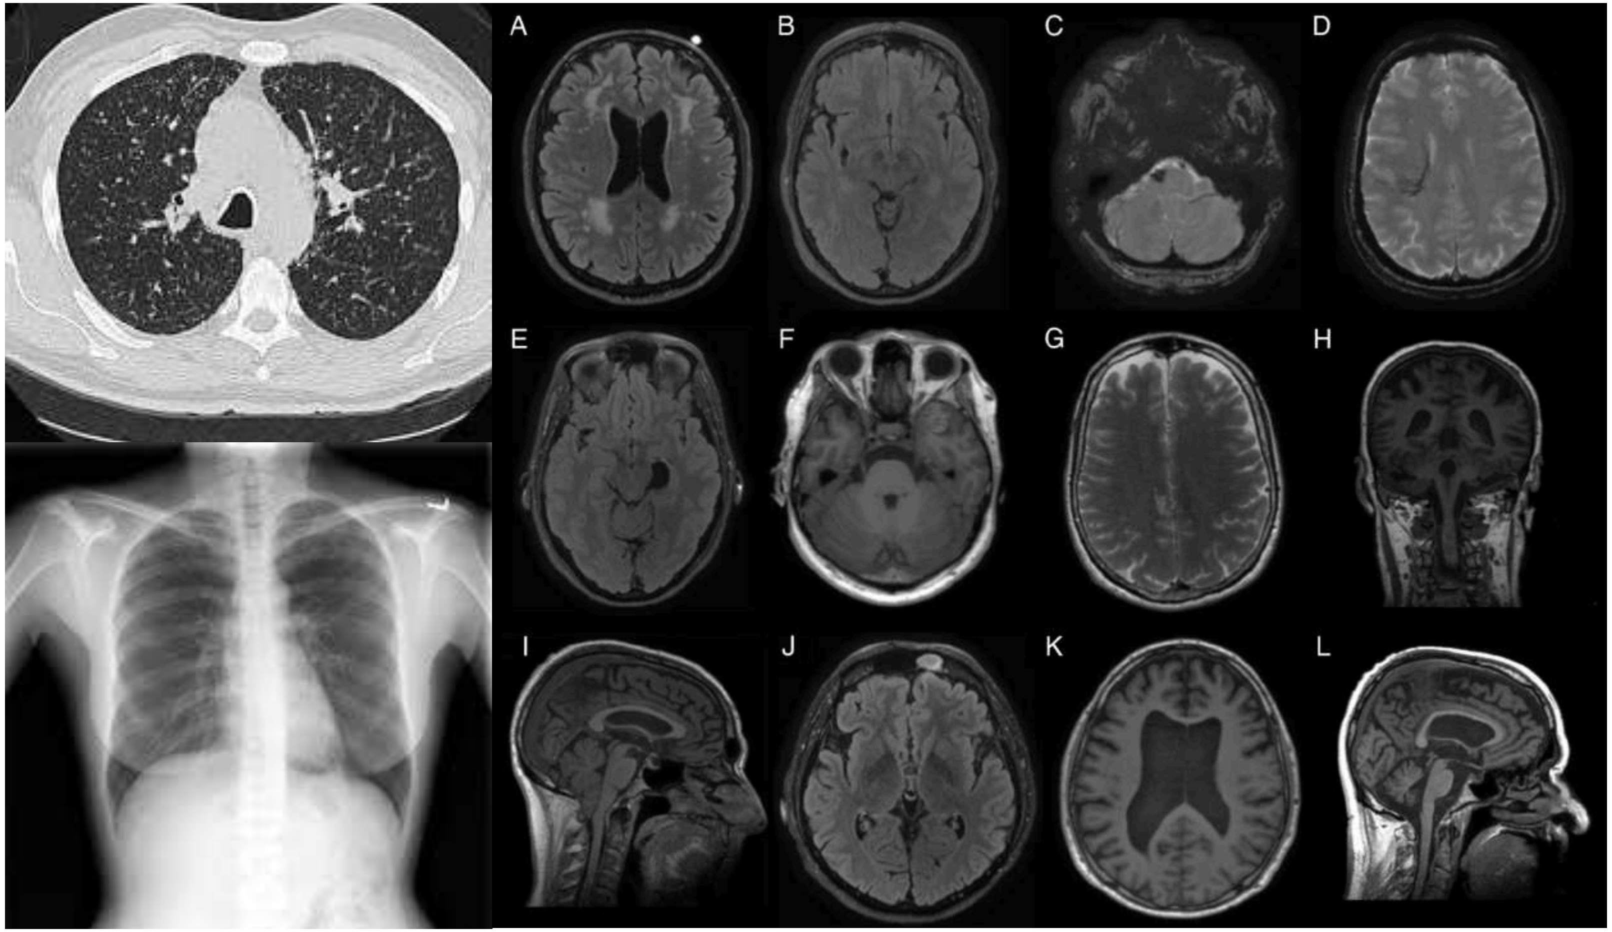

Zhang, et al. "Colorful image colorization." In European Conference on Computer Vision (2016).

Self-supervised learning

To learn representation directly from data itself

Example #4: image context prediction

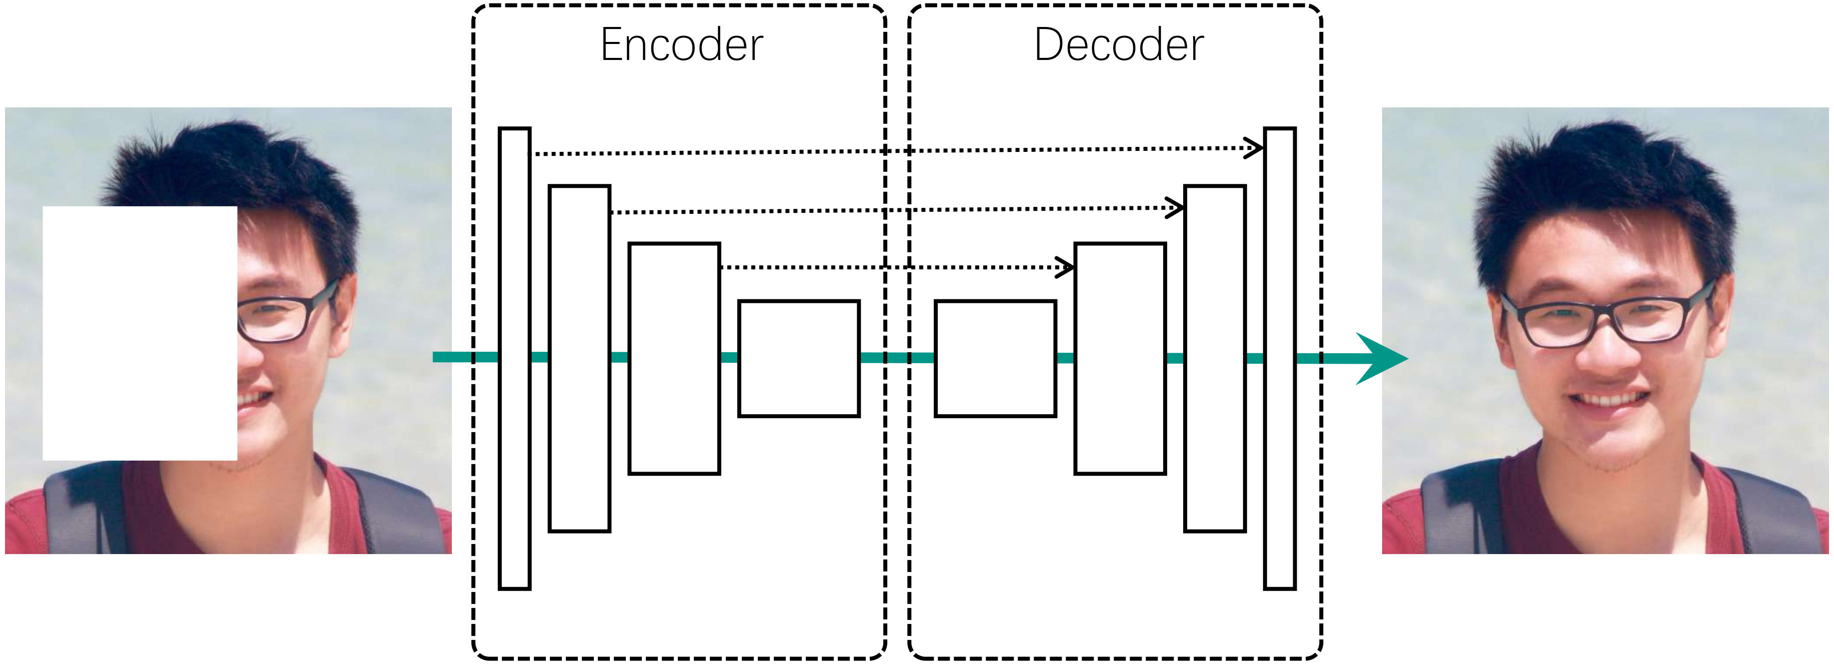

Pathak et al. "Context encoders: Feature learning by inpainting." In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2016).

Summary: Self-supervised learning

- Learn directly from data itself
- Design the input-output pairs
- Predict from the disrupted original data

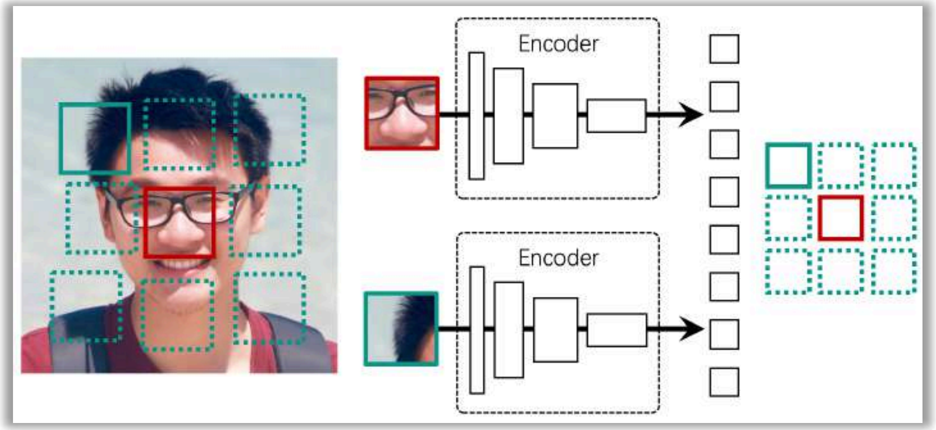

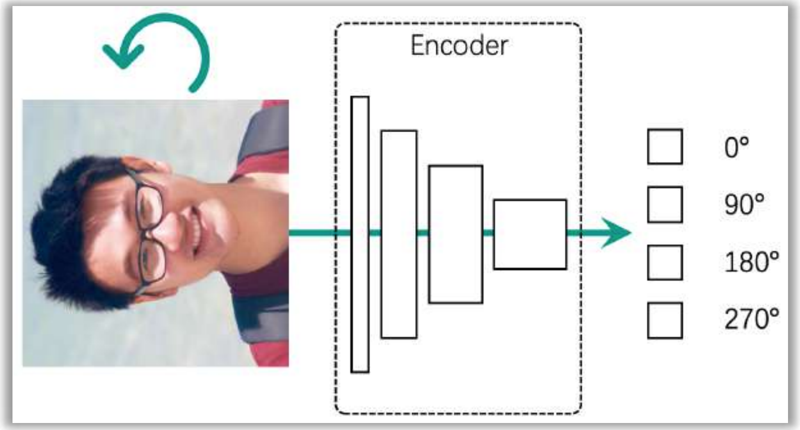

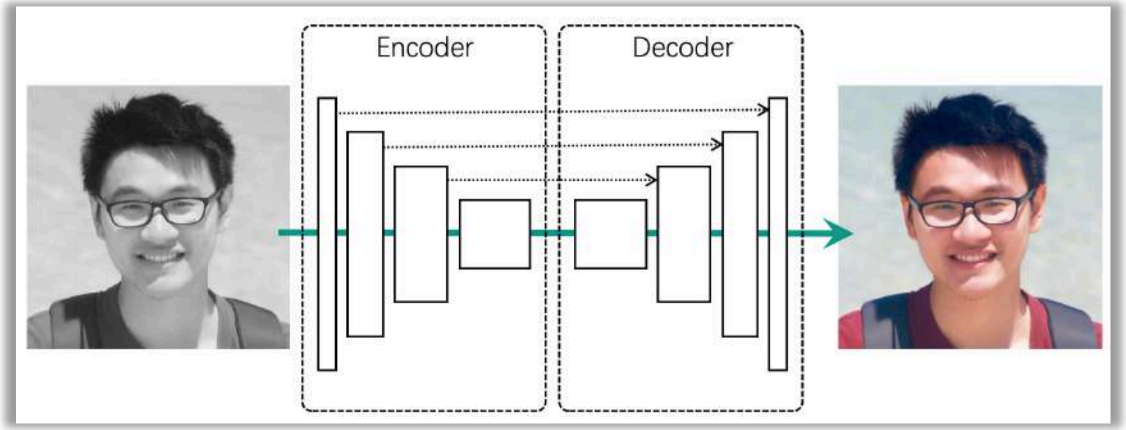

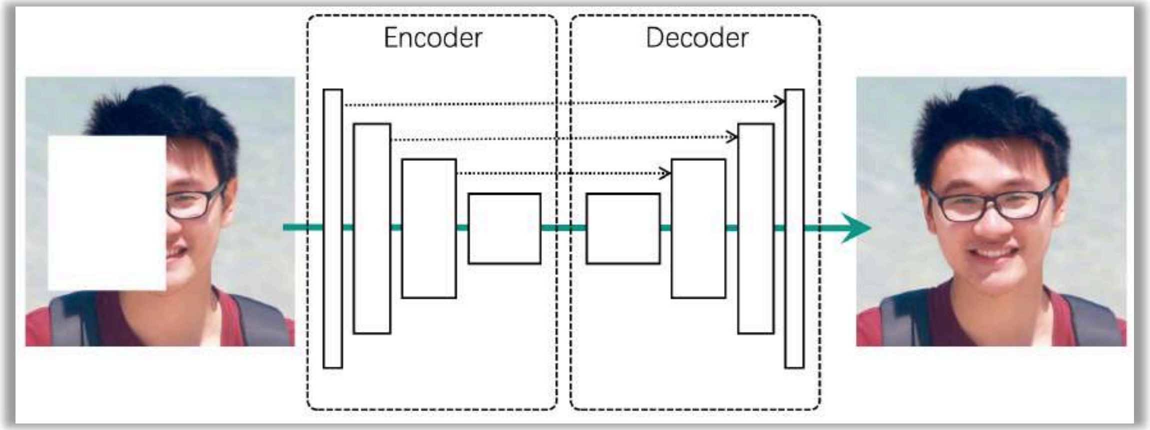

Example #5: exemplar learning

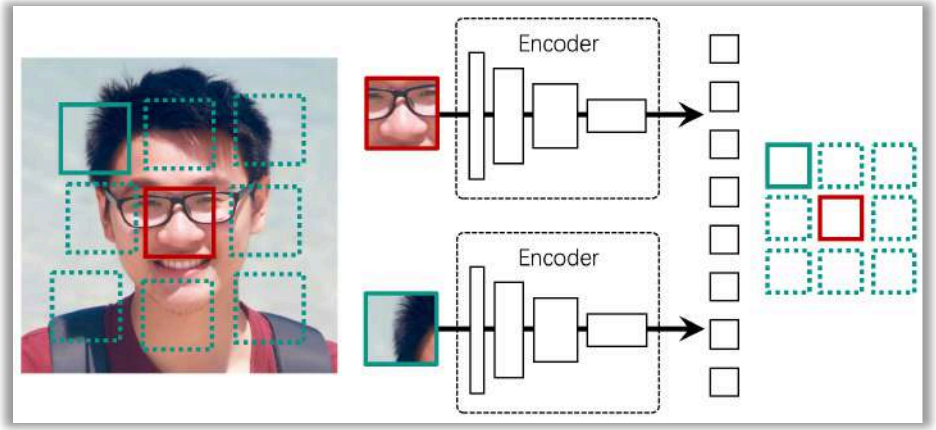

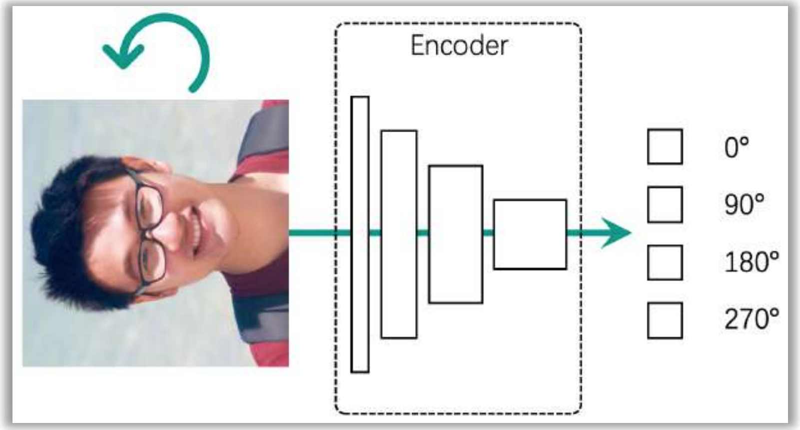

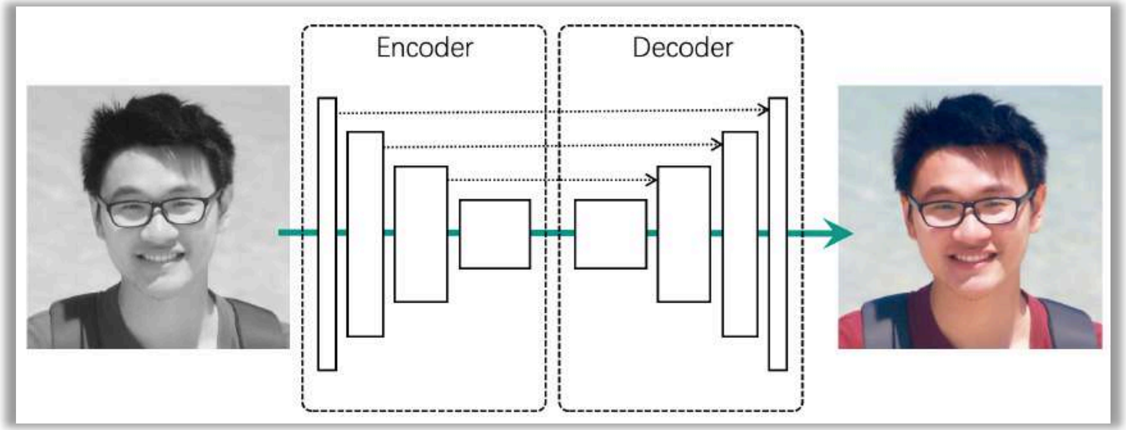

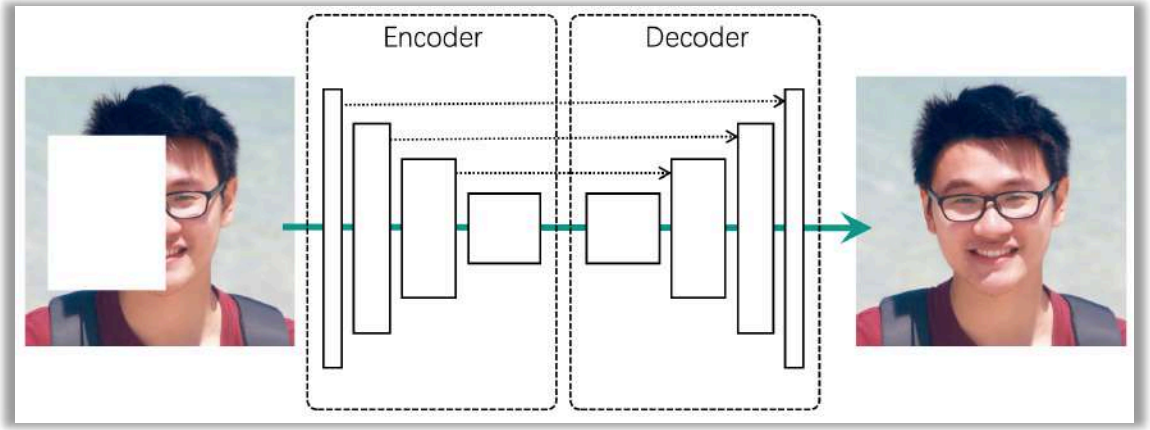

We design it as a simple image restoration task, through which, the model can learn representation directly from image data itself.

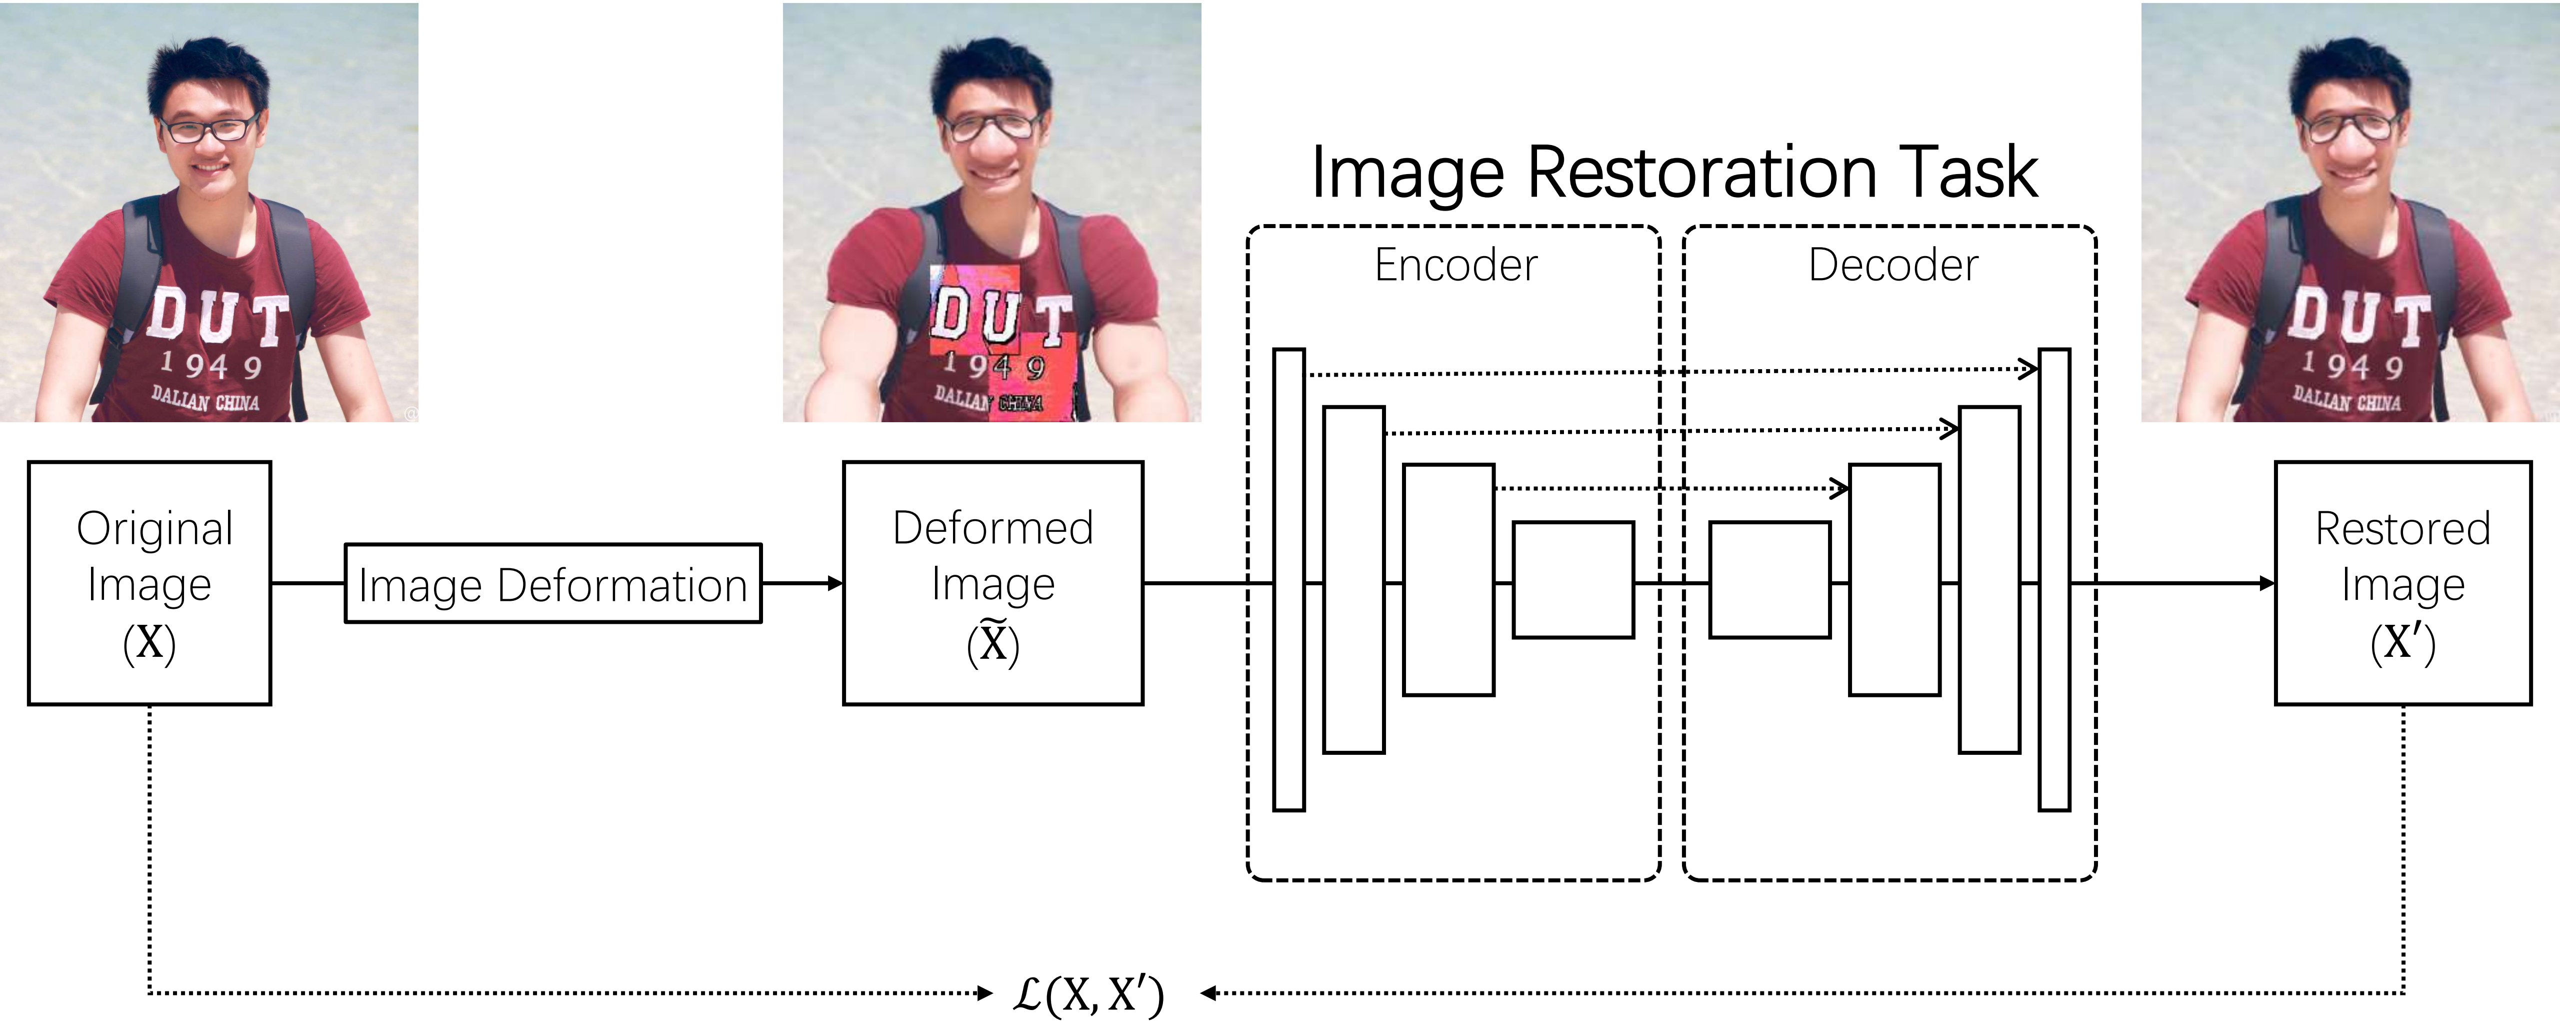

Our self-supervised learning framework is scalable because it is easy to incorporate any other meaningful image deformations.

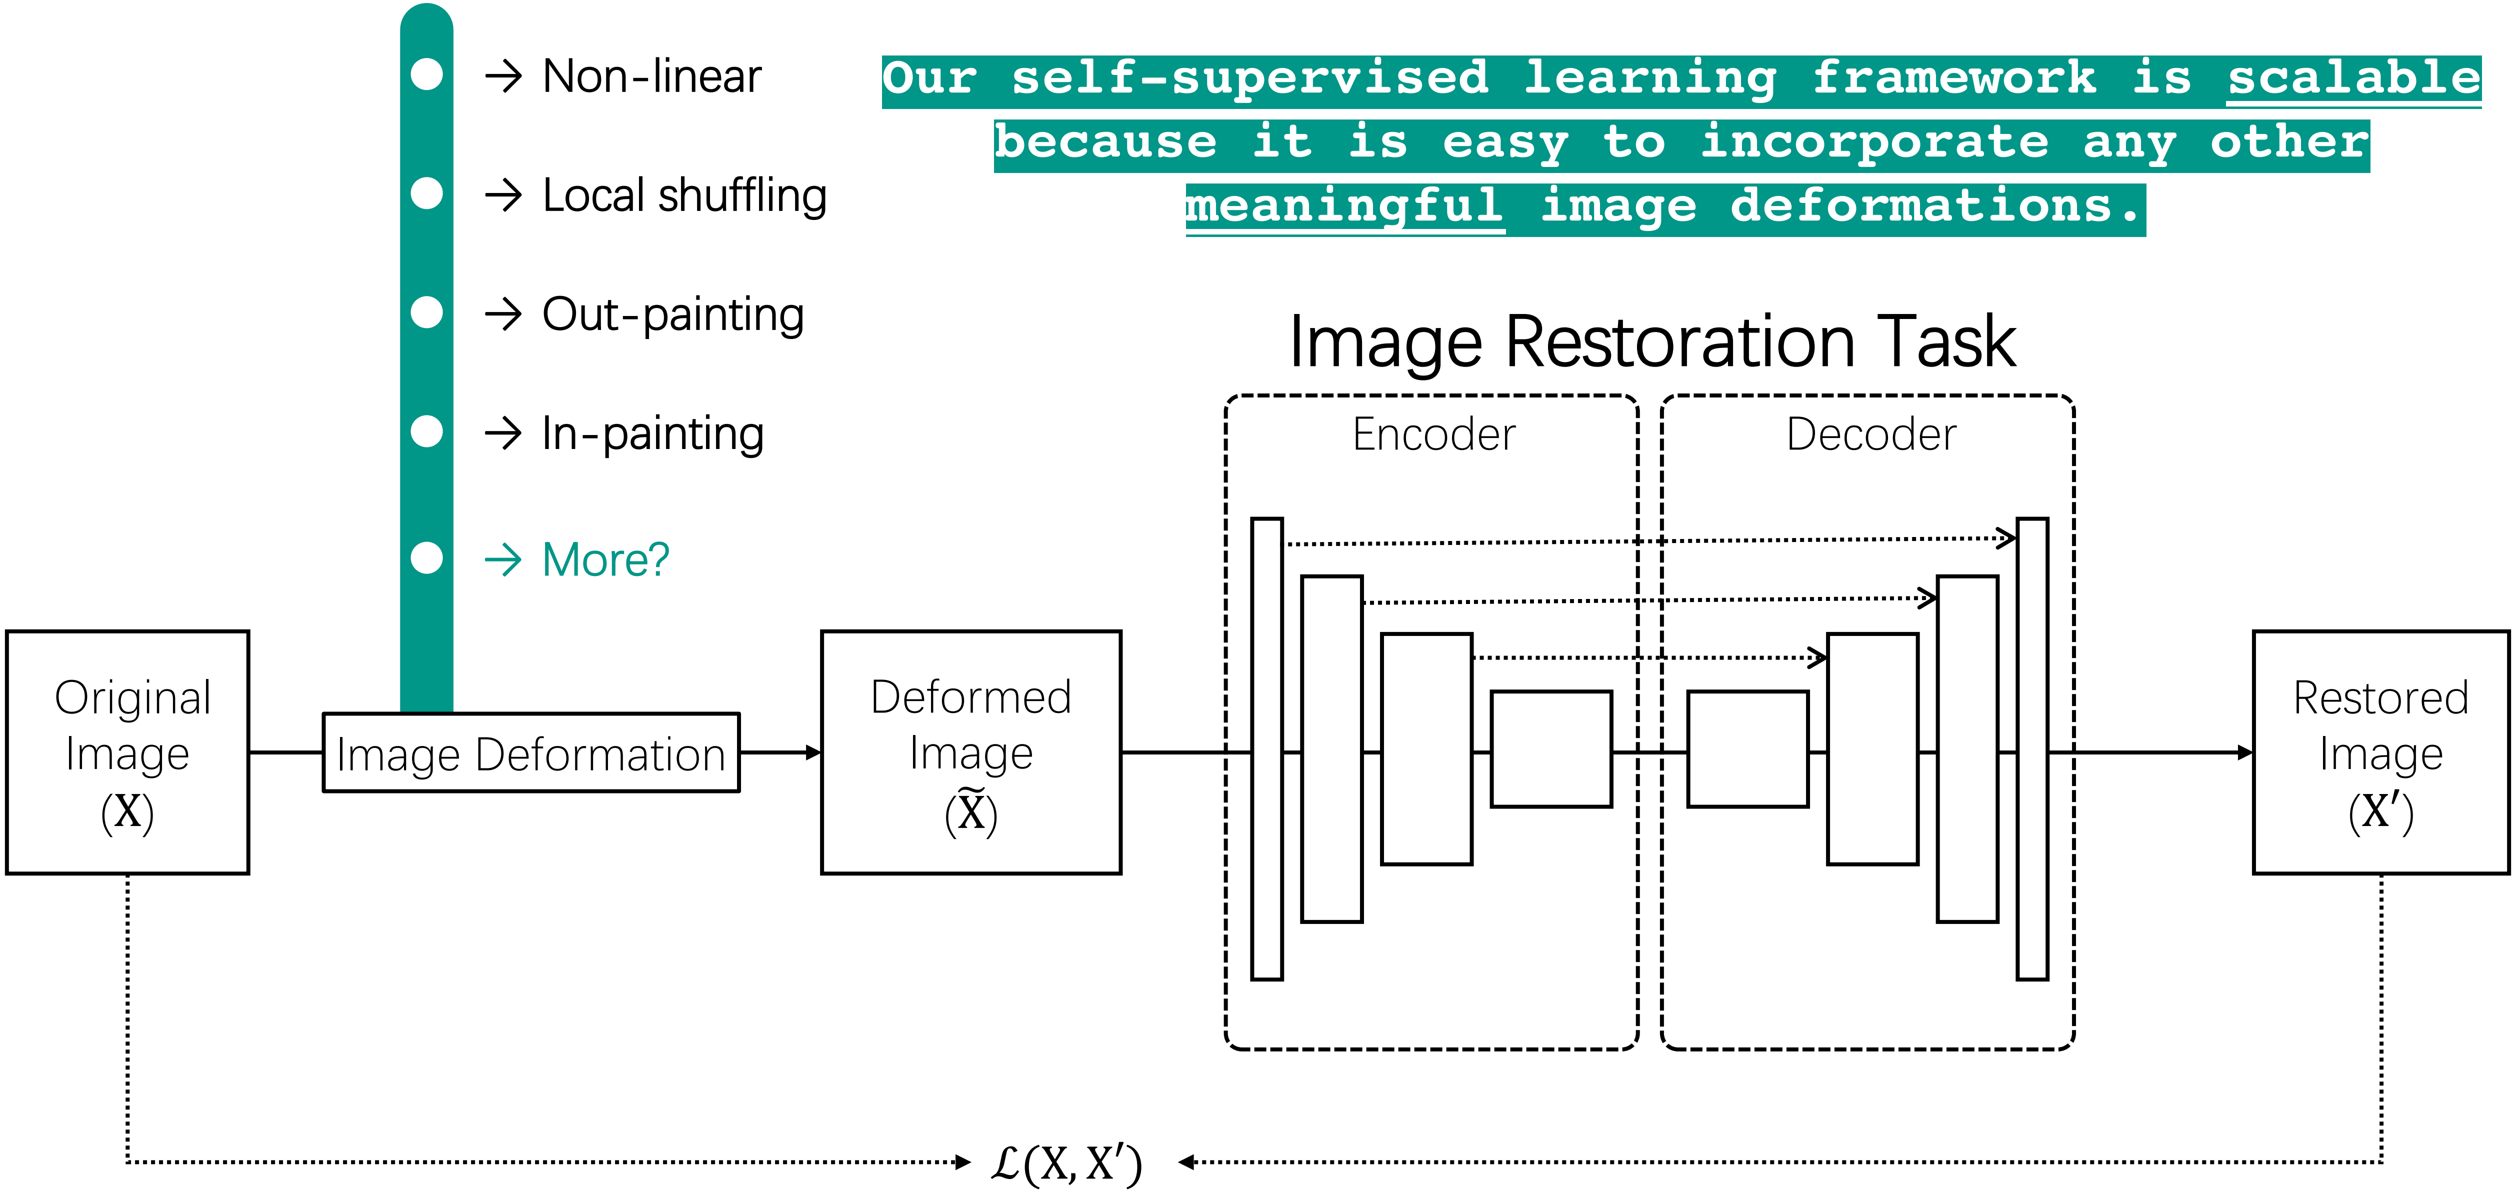

Image deformation vs. data augmentation?

- https://github.com/albu/albumentations

Example	Image deformation	Data augmentation
Translation	✗	✓

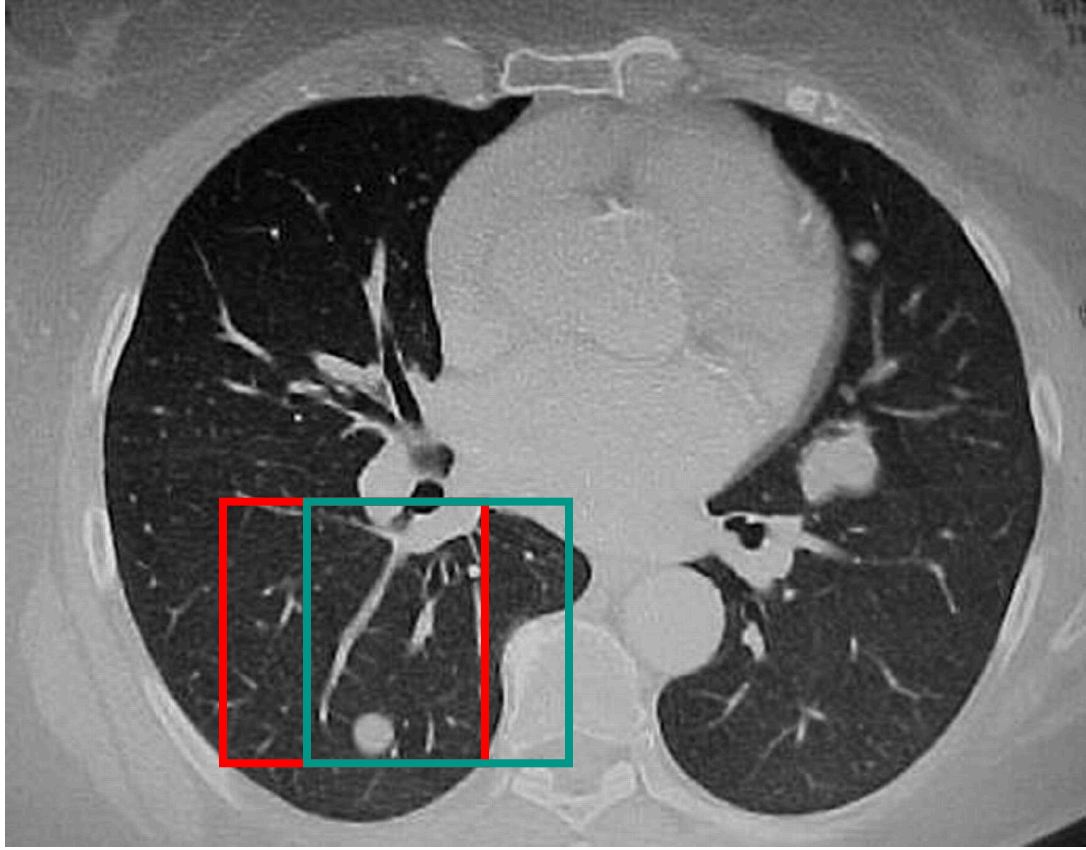

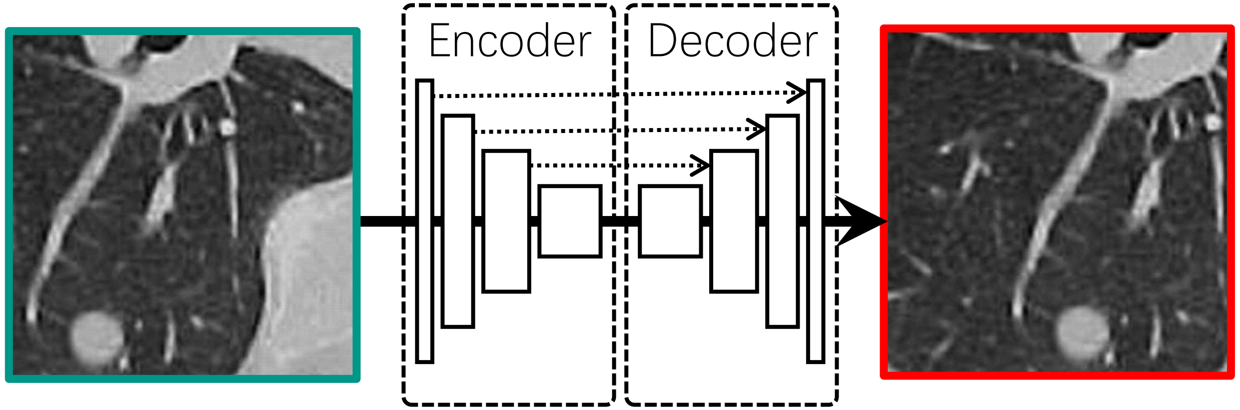

Image deformation vs. data augmentation?

- https://github.com/albu/albumentations

Example	Image deformation	Data augmentation
Translation	✗	✓
Rotation	✗	✓

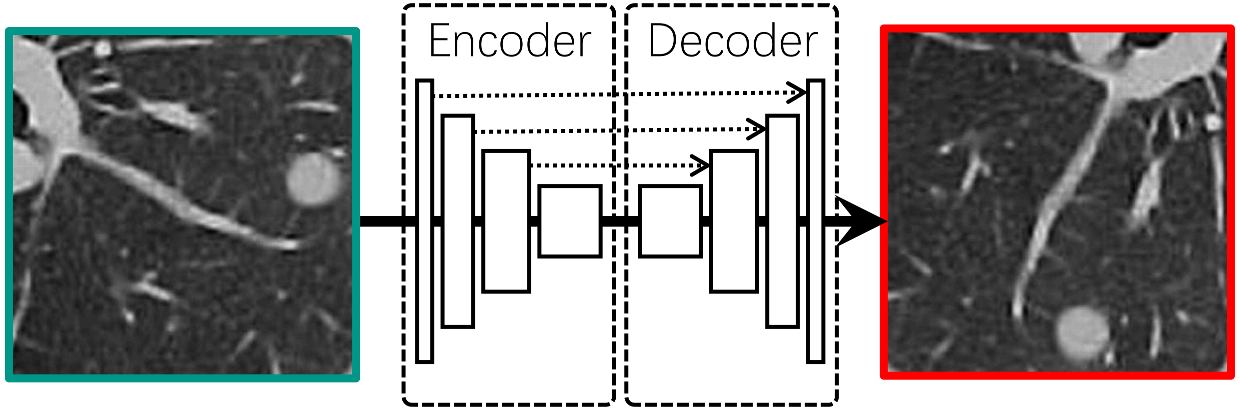

Image deformation vs. data augmentation?

- https://github.com/albu/albumentations

Example	Image deformation	Data augmentation
Translation	✗	✓
Rotation	✗	✓
Flipping	✗	✓

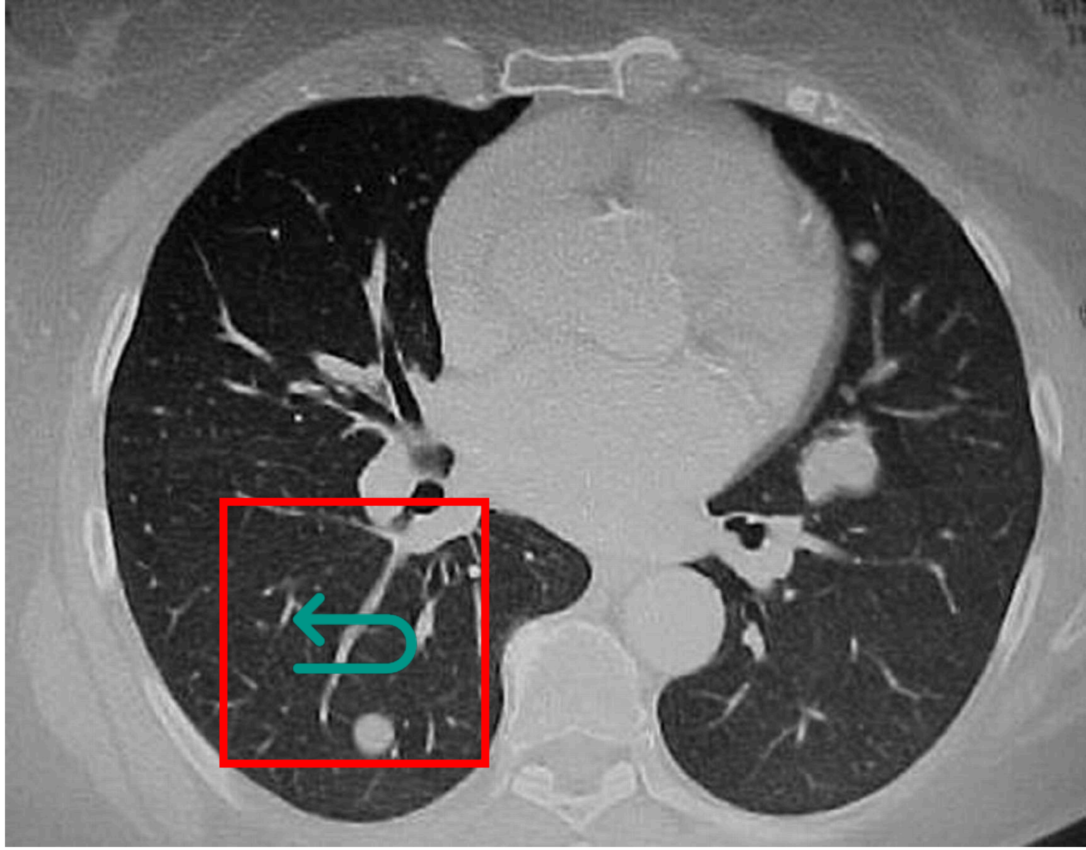

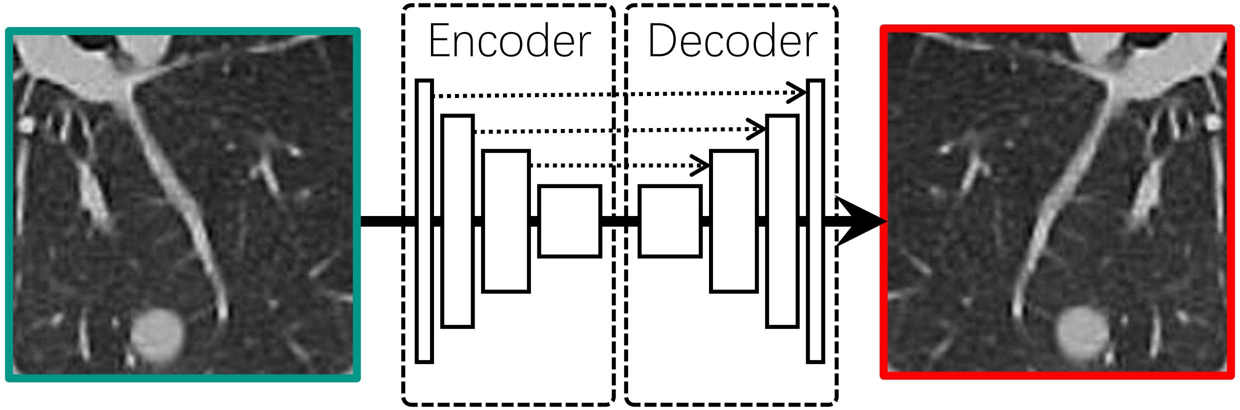

Image deformation vs. data augmentation?

- https://github.com/albu/albumentations

Example	Image deformation	Data augmentation
Translation	✗	✓
Rotation	✗	✓
Flipping	✗	✓
Scaling	✗	✓

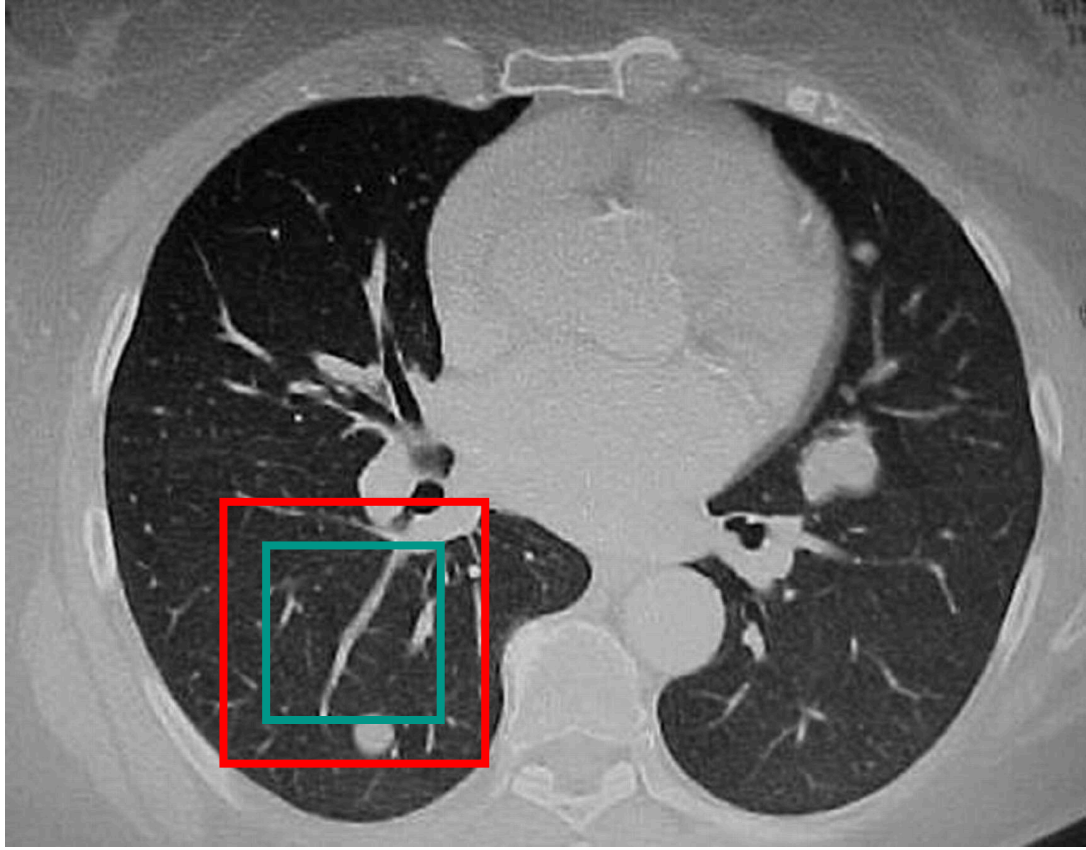

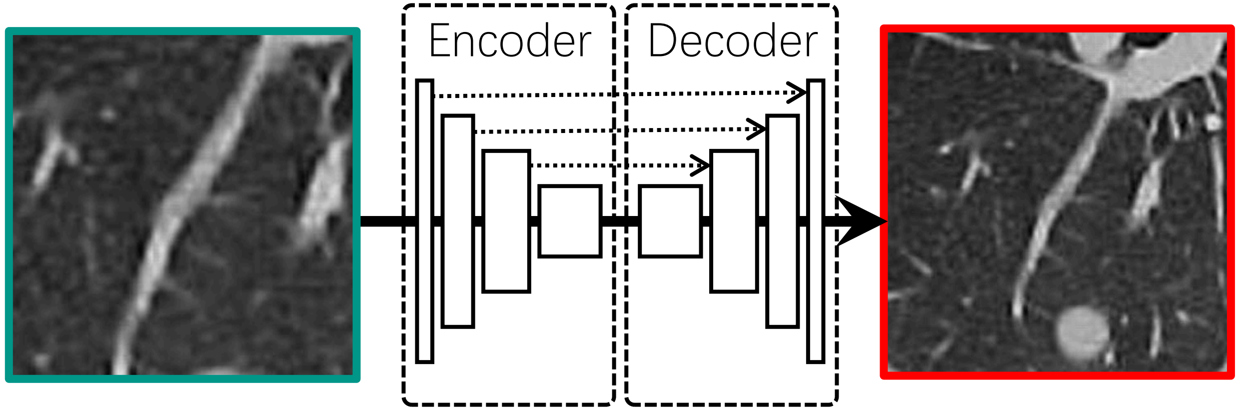

Image deformation vs. data augmentation?

- https://github.com/albu/albumentations

Example	Image deformation	Data augmentation
Translation	✗	✓
Rotation	✗	✓
Flipping	✗	✓
Scaling	✗	✓
Blur	✓	✓

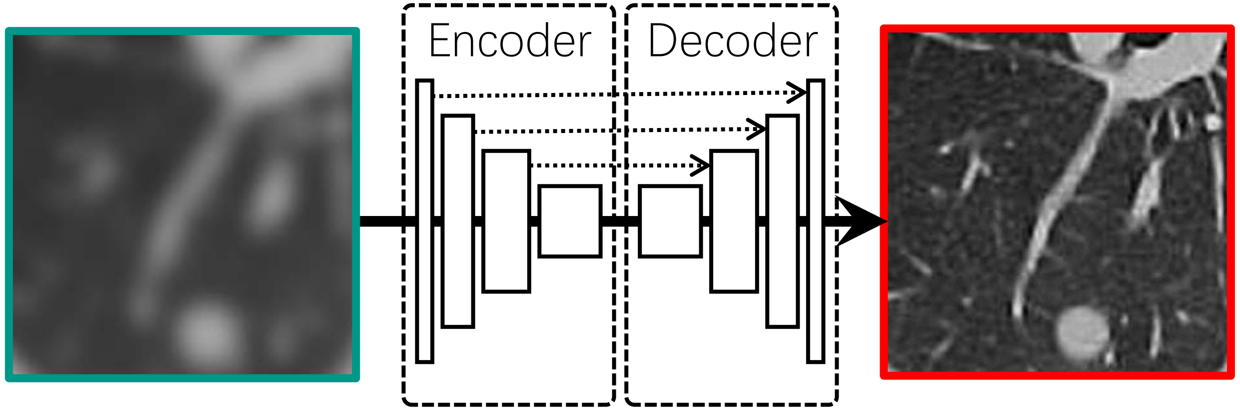

Image deformation vs. data augmentation?

- https://github.com/albu/albumentations

Example	Image deformation	Data augmentation
Translation	✗	✓
Rotation	✗	✓
Flipping	✗	✓
Scaling	✗	✓
Blur	✓	✓
Noise	✓	✓

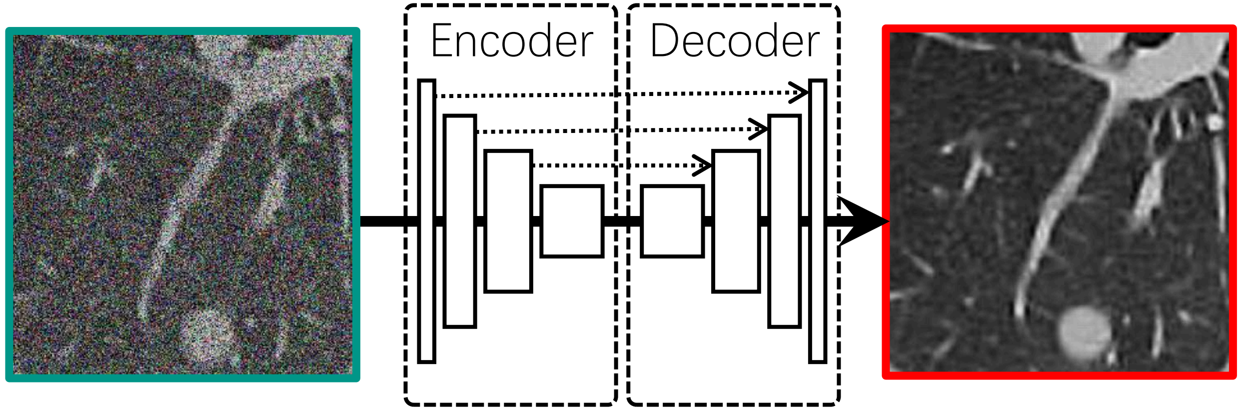

Image deformation vs. data augmentation?

- https://github.com/albu/albumentations

Example	Image deformation	Data augmentation
Translation	✗	✓
Rotation	✗	✓
Flipping	✗	✓
Scaling	✗	✓
Blur	✓	✓
Noise	✓	✓
…	?	?

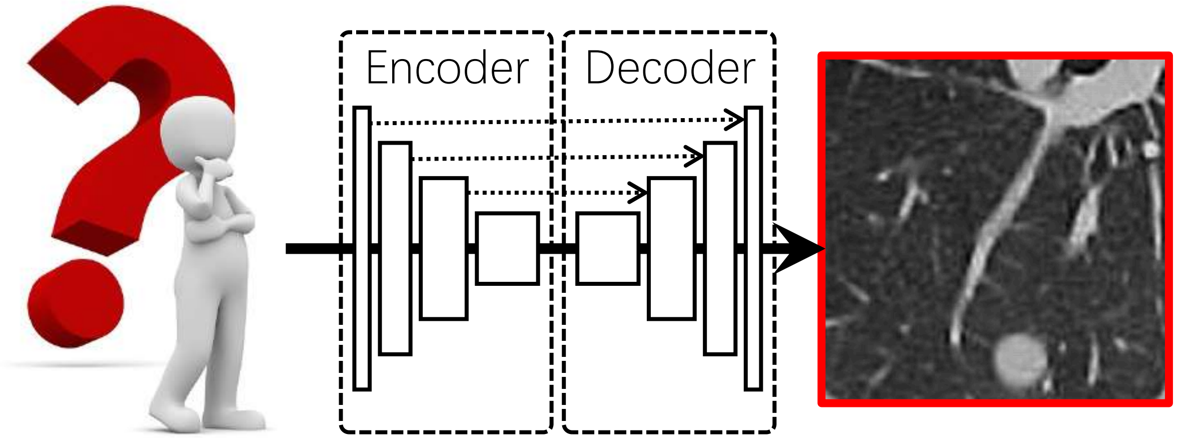

Image deformation vs. data augmentation?

- https://github.com/albu/albumentations

To incorporate other meaningful image deformations into our framework, the deformation should belong to pixel-level transform, rather than spatial-level transform.

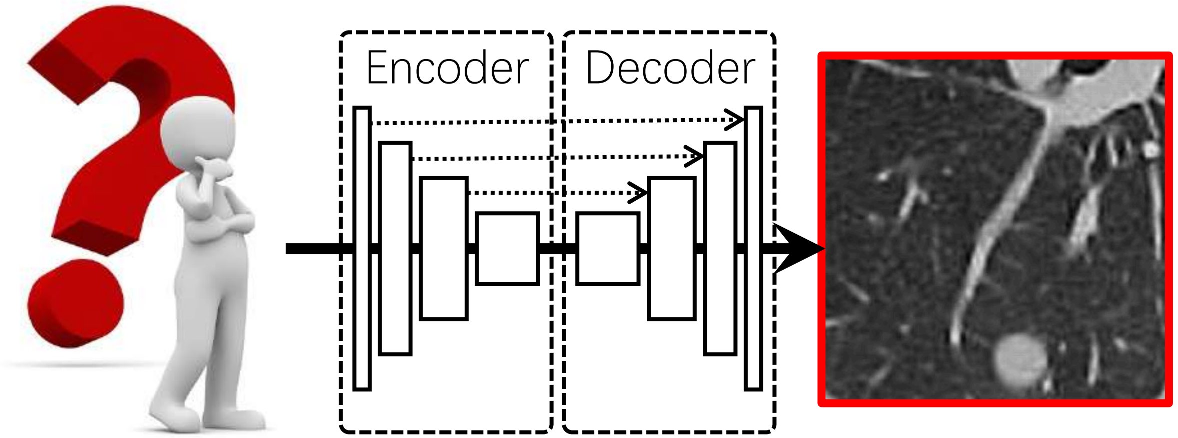

Image deformation vs. data augmentation?

- https://github.com/albu/albumentations

Why the proposed image deformations in your paper work?

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.

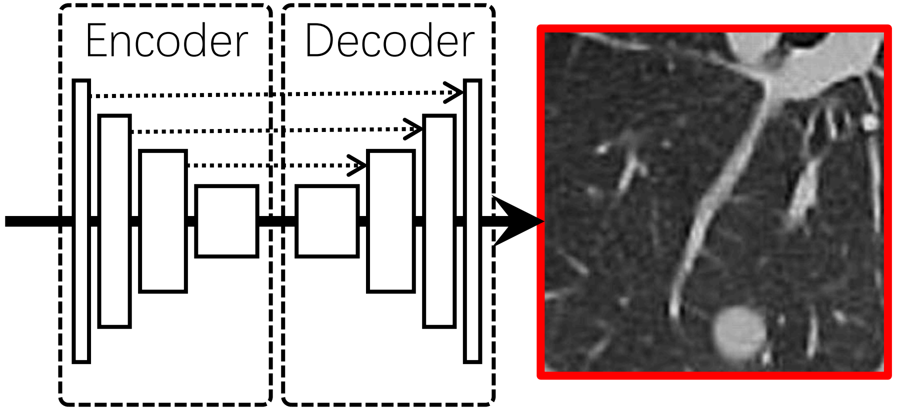

Medical ImageNet?

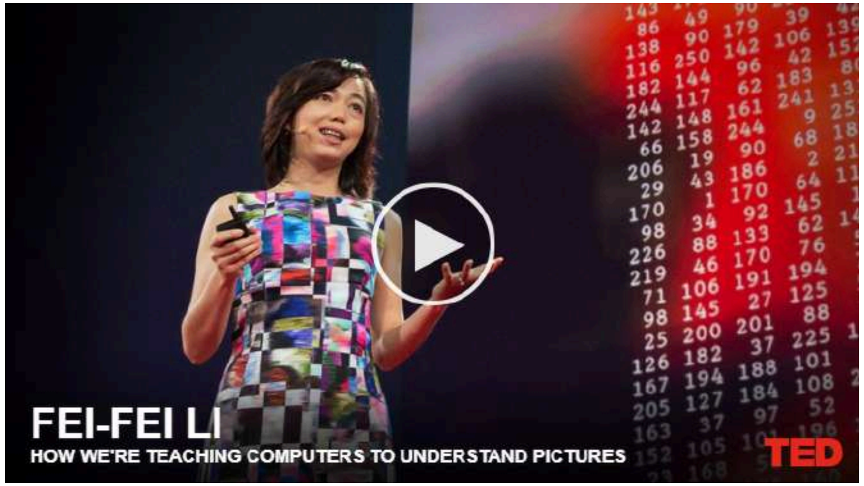

IMAGENET

14,197,122 images, 21841 synsets indexed

Explore Download Challenges Publications Updates About

Not logged in. Login I Signup

ImageNet is an image database organized according to the WordNet hierarchy (currently only the nouns), in which each node of the hierarchy is depicted by hundreds and thousands of images. Currently we have an average of over five hundred images per node. We hope ImageNet will become a useful resource for researchers, educators, students and all of you who share our passion for pictures.

Click here to learn more about ImageNet, Click here to join the ImageNet mailing list.

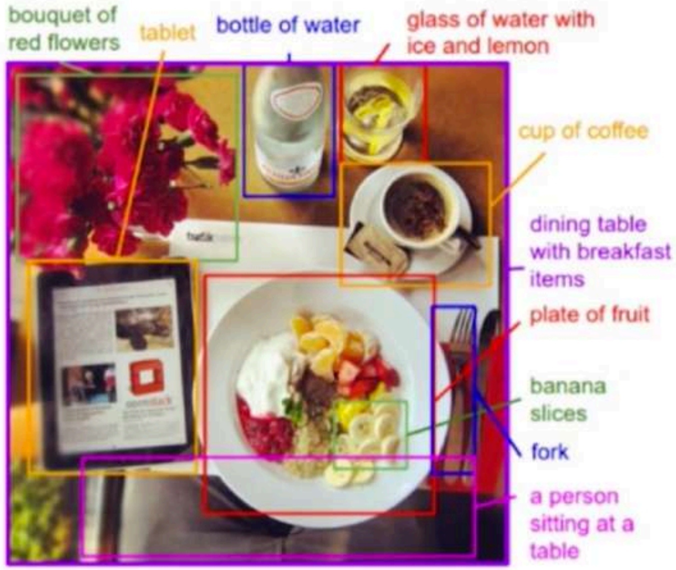

Karpathy, Andrej & Li, Fei Fei. Deep Visual-Semantic Alignments for Generating Image Descriptions, CVPR, 2015

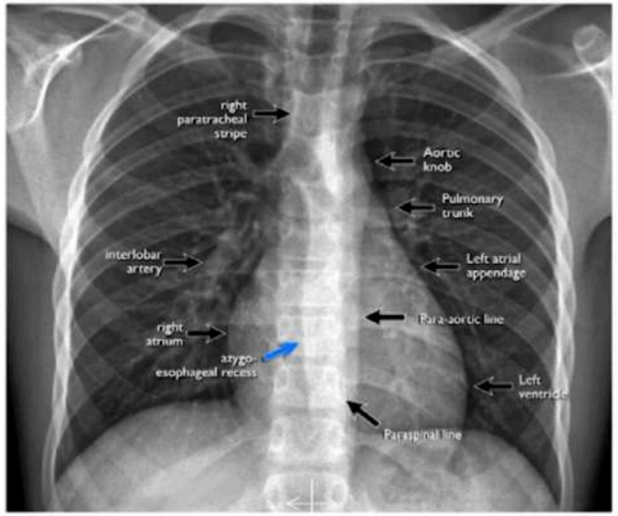

http://www.radiologyassistant.nl/

"Medical ImageNet"*

A cloud-based, petabyte-scale, searchable, repository of diagnostic imaging studies for developing intelligent image analysis systems

*Thanks to Fei Fei Li

Stanford | MEDICINE

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

This result is unprecedented because

no self-supervised methods have thus far performed as well as Models ImageNet.

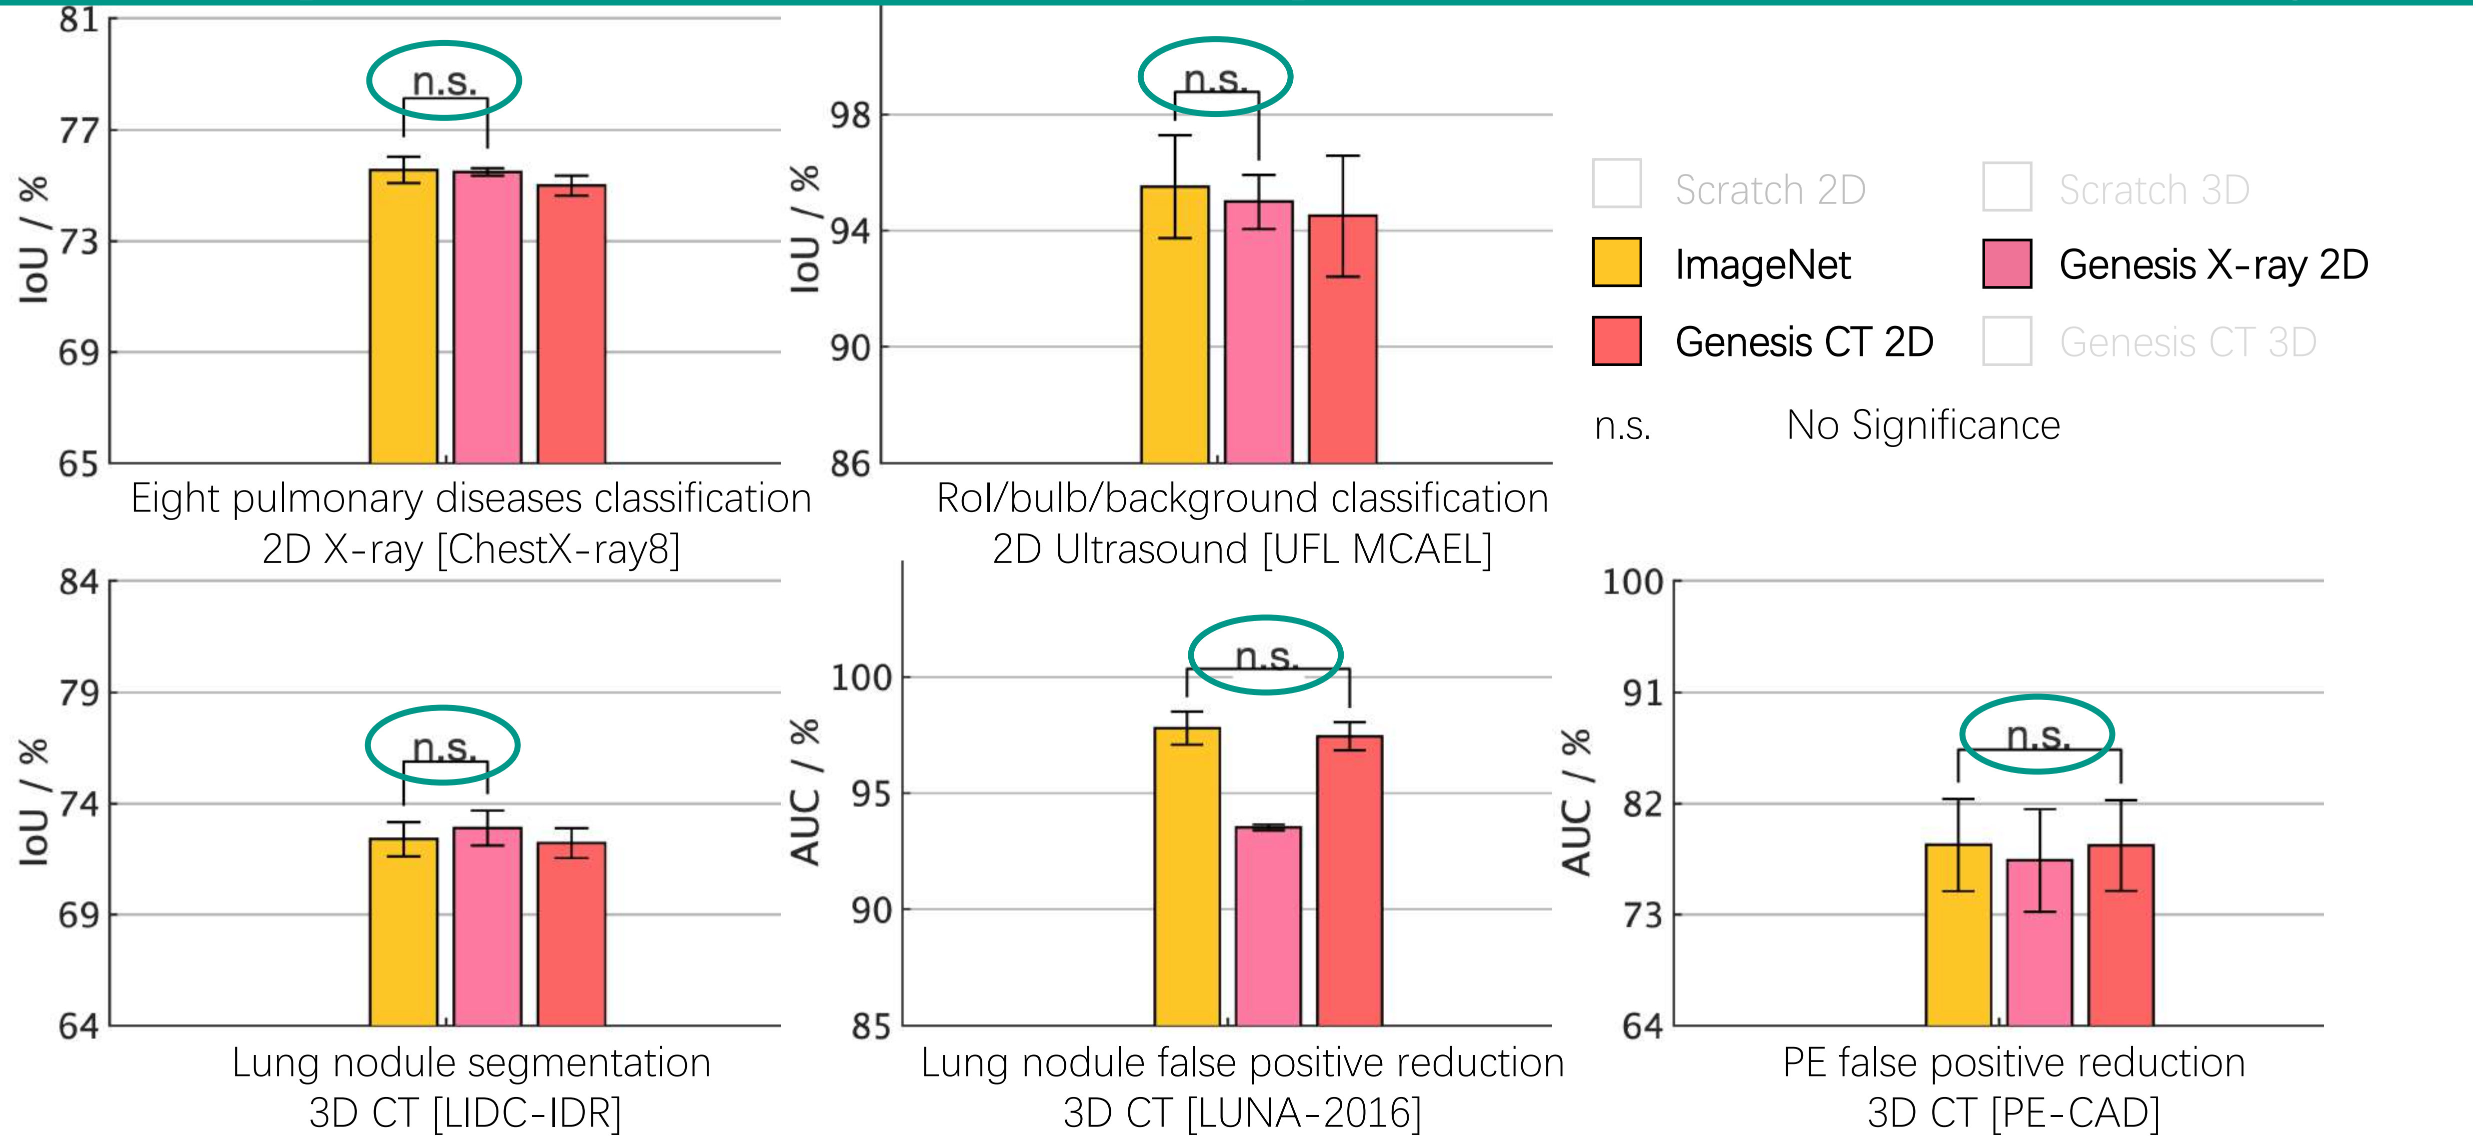

Family		ImageNet		Places205	
		Prev.	Ours	Prev.	Ours
A	Rotation[11]	38.7	**55.4**	35.1	**48.0**
R	Exemplar[8]	31.5	46.0	-	42.7
R	Rel. Patch Loc.[8]	36.2	51.4	-	45.3
A	Jigsaw[34, 51]	34.7	44.6	35.5	42.2
V	CC+vgg-Jigsaw++[36]	37.3	-	37.5	-
A	Counting[35]	34.3	-	36.3	-
A	Split-Brain[51]	35.4	-	34.1	-
V	DeepClustering[3]	**41.0**	-	**39.8**	-
R	CPC[37]	48.7†	-	-	-
R	Supervised RevNet50	74.8	74.4	-	58.9
R	Supervised ResNet50 v2	76.0	75.8	-	61.6
V	Supervised VGG19	72.7	75.0	58.9	61.5

† *marks results reported in unpublished manuscripts.*

Kolesnikov et al. "Revisiting Self-Supervised Visual Representation Learning." In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019). https://github.com/google/revisiting-self-supervised

Method	Pretext Tasks	Classification	Detection	Segmentation
ImageNet Labels [8]	—	79.9	56.8	48.0
Random(Scratch) [8]	—	57.0	44.5	30.1
ContextEncoder [19]	Generation	56.5	44.5	29.7
BiGAN [122]	Generation	60.1	46.9	35.2
ColorfulColorization [18]	Generation	65.9	46.9	35.6
SplitBrain [42]	Generation	67.1	46.7	36.0
RankVideo [38]	Context	63.1	47.2	$35.4^{\dagger}$
PredictNoise [46]	Context	65.3	49.4	$37.1^{\dagger}$
JigsawPuzzle [20]	Context	67.6	53.2	37.6
ContextPrediction [41]	Context	65.3	51.1	—
Learning2Count [130]	Context	67.7	51.4	36.6
DeepClustering [44]	**Context**	**73.7**	**55.4**	**45.1**
WatchingVideo [81]	Free Semantic Label	61.0	52.2	—
CrossDomain [30]	Free Semantic Label	68.0	52.6	—
AmbientSound [154]	Cross Modal	61.3	—	—
TiedToEgoMotion [95]	Cross Modal	—	41.7	—
EgoMotion [94]	Cross Modal	54.2	43.9	—

Jing, et al. “Self-supervised Visual Feature Learning with Deep Neural Networks: A Survey." arXiv preprint arXiv:1902.06162 (2019).
https://arxiv.org/pdf/1902.06162.pdf

Medical ImageNet?

Models Genesis are not designed to replace such a large, strongly annotated dataset for medical image analysis like ImageNet for computer vision, but rather helping create one.

Models Medical ImageNet > Models Genesis?

1. Millions of systematic annotated medical images
2. Disease/organ class imbalance
3. Pixel/voxel utilization rate
4. Availability of medical images

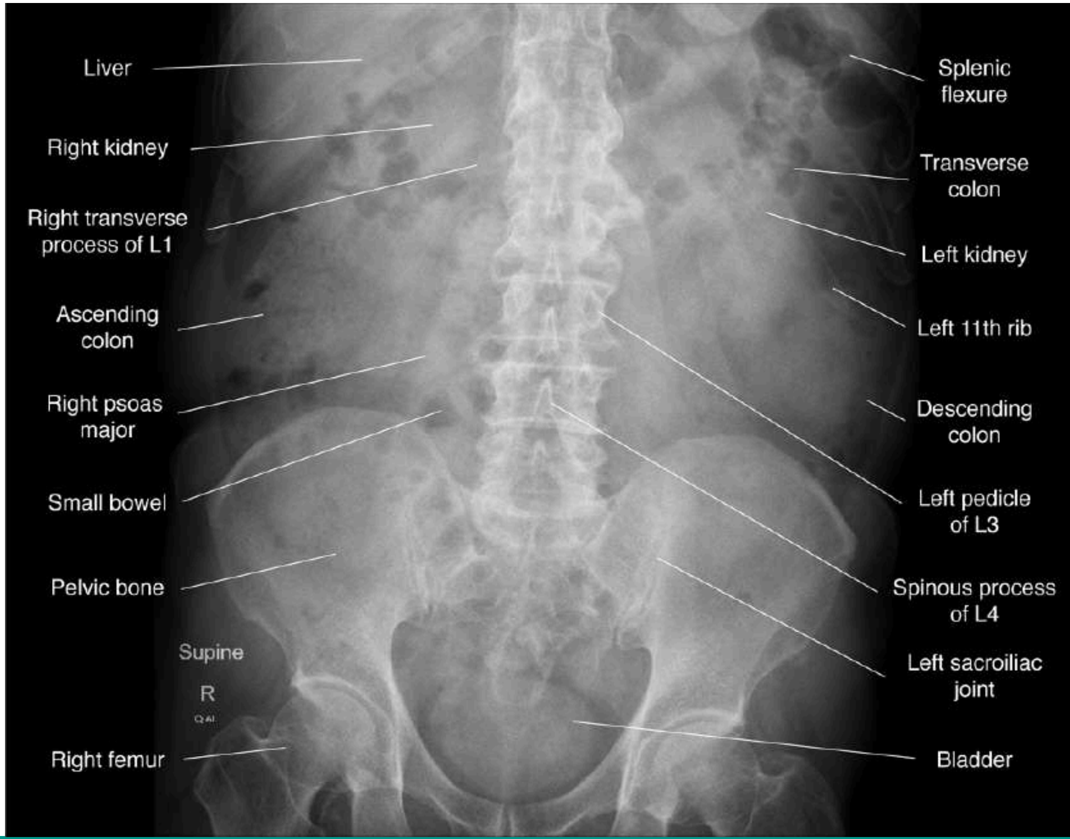

We hope that our collective efforts will lead to the Holy Grail of Models Genesis, effective across diseases, organs, and modalities.

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

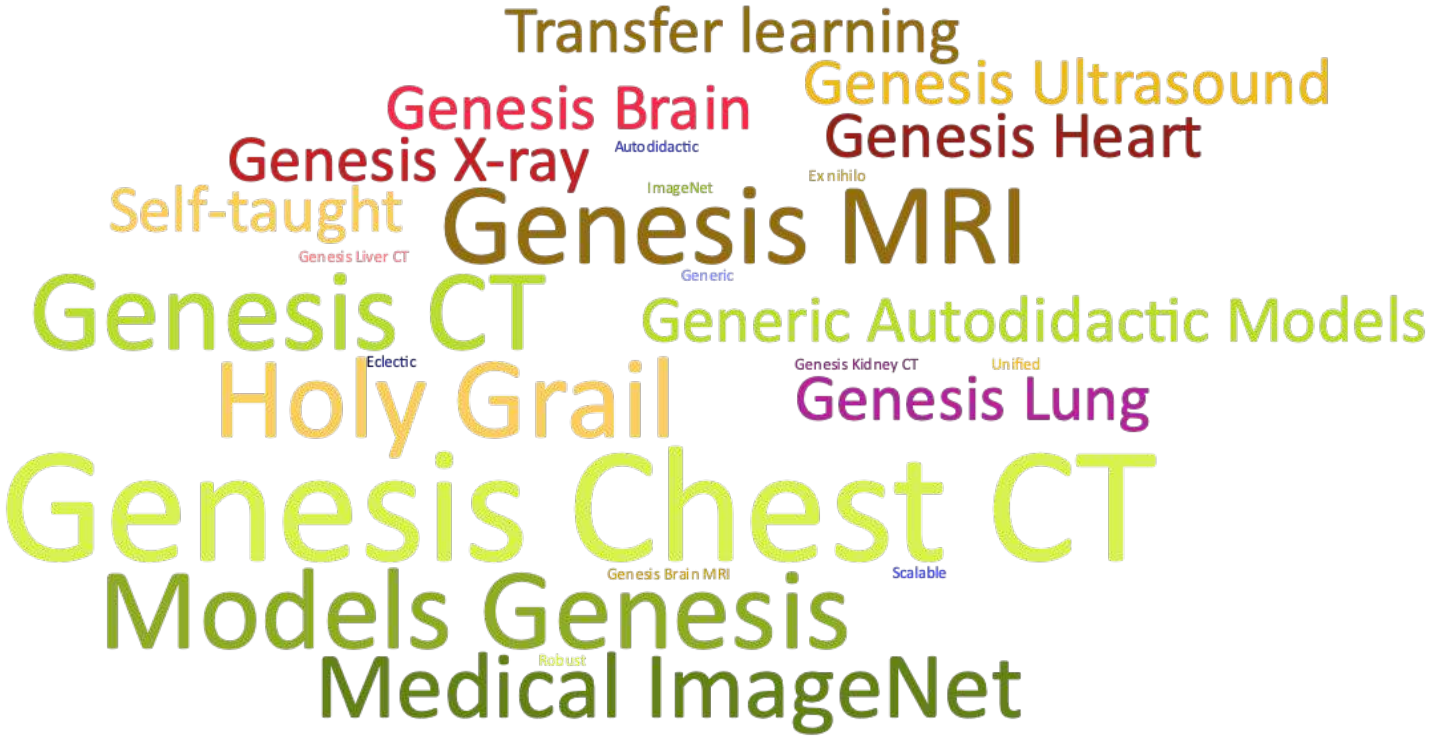

Paper

This repository provides the official Keras 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)

Hands-on

https://github.com/MrGiovanni/ModelsGenesis

Fine-tune from our pre-trained Models Genesis

- Download Genesis_Chest_CT.h5 from Google Drive or Baidu Wangpan

Thank you!
Please download the pre-trained models via the link below

https://drive.google.com/file/d/11yGsC8LL9WKO47vCWeU0Axg0yHmbvCk_/view?usp=sharing

Alternatively, for the one who cannot access Google, please download here
Link: https://pan.baidu.com/s/1qXT5XQ4KoQC4LXUltR8sdA Pass: rvr9

Please contact Zongwei Zhou if you have problem downloading the pre-trained models at
zongweiz@asu.edu

powered by www.wjx.cn™

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

Fine-tune from our pre-trained Models Genesis

- Download Genesis_Chest_CT.h5 from Google Drive or Baidu Wangpan
- Ours is the best? Not necessary
- Input size? Divisible by 16 but (N, 1, 64, 64, 32) is preferred

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

Fine-tune from our pre-trained Models Genesis

- Download Genesis_Chest_CT.h5 from Google Drive or Baidu Wangpan
- Ours is the best? Not necessary
- Input size? Divisible by 16 but (N, 1, 64, 64, 32) is preferred

Build your own Models Genesis

- Generate cube pairs using infinite_generator_3D.py
- Pre-process the data
- Self-supervised learning using Genesis_Chest_CT.py

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

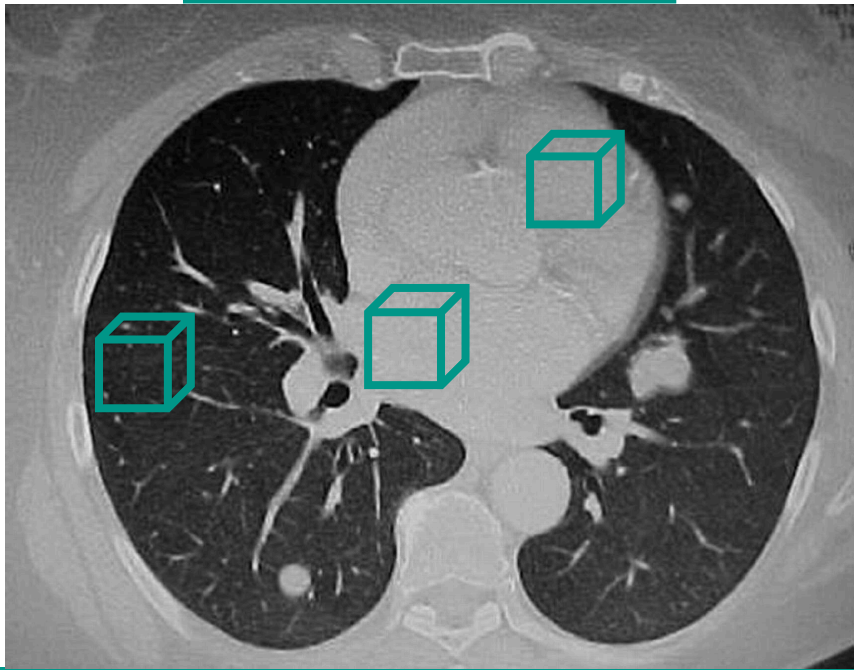

AI研习社

Fine-tune from our pre-trained Models Genesis

- Download Genesis_Chest_CT.h5 from Google Drive or Baidu Wangpan
- Ours is the best? Not necessary
- Input size? Divisible by 16 but (N, 1, 64, 64, 32) is preferred

Build your own Models Genesis

- Generate cube pairs using infinite_generator_3D.py
- Pre-process the data
- Self-supervised learning using Genesis_Chest_CT.py

Preferred ☺

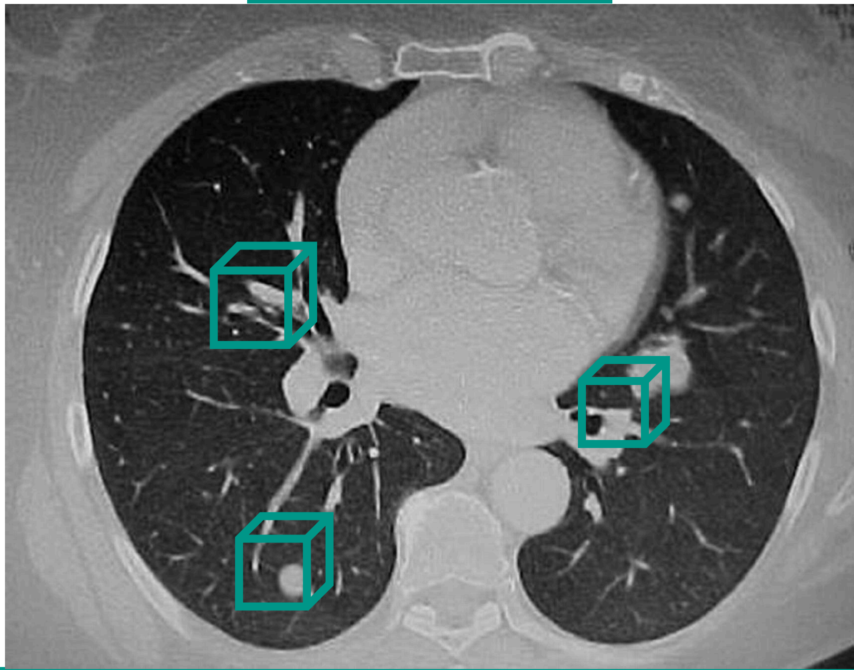

We make the development of Models Genesis open science and invite researchers around the world to contribute to this effort.

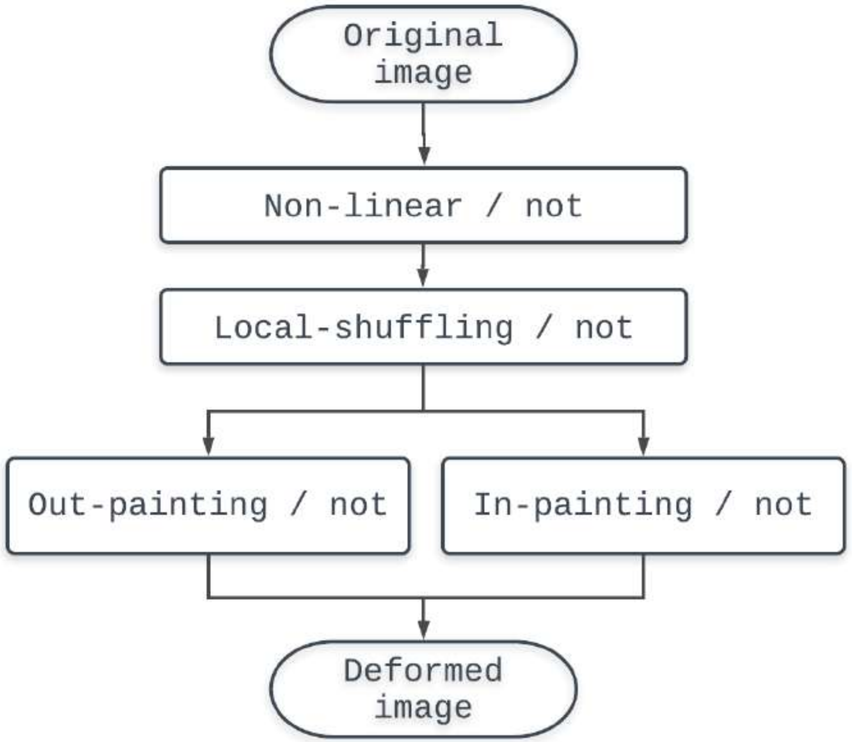

Deformation	1	2	3	4	5	6	7	8	9	10	11	12
Non-linear		✓	✓	✓	✓	✓	✓					
Local-shuffling			✓			✓	✓	✓	✓	✓		
Out-painting				✓		✓			✓		✓	
In-painting					✓		✓			✓		✓

- How to combine different image deformations?
  - Each deformation is independently applied to a patch with a predefined probability, while out-painting and in-painting are considered mutually exclusive.

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

AI研习社

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