



Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Models Genesis

Zongwei Zhou^a, Vatsal Sodha^b, Jiaxuan Pang^b, Michael B. Gotway^c, Jianming Liang^{a,*}^a Department of Biomedical Informatics, Arizona State University, Scottsdale, AZ 85259, USA^b School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe, AZ 85281 USA^c Department of Radiology, Mayo Clinic, Scottsdale, AZ 85259, USA

ARTICLE INFO

Article history:

Received 10 March 2020

Revised 12 August 2020

Accepted 14 September 2020

Available online 13 October 2020

MSC:

41A05

41A10

65D05

65D17

Keywords:

3D Deep learning

Representation learning

Transfer learning

Self-supervised learning

ABSTRACT

Transfer learning from *natural* images to *medical* images has been established as one of the most practical paradigms in deep learning for medical image analysis. To fit this paradigm, however, 3D imaging tasks in the most prominent imaging modalities (e.g., CT and MRI) have to be reformulated and solved in 2D, losing rich 3D anatomical information, thereby inevitably compromising its performance. To overcome this limitation, we have built a set of models, called Generic Autodidactic Models, nicknamed Models Genesis, because they are created *ex nihilo* (with no manual labeling), self-taught (learnt by self-supervision), and generic (served as source models for generating application-specific target models). Our extensive experiments demonstrate that our Models Genesis significantly outperform learning from scratch and existing pre-trained 3D models in all five target 3D applications covering both segmentation and classification. More importantly, learning a model from scratch *simply in 3D* may not necessarily yield performance better than transfer learning from ImageNet in 2D, but our Models Genesis consistently top any 2D/2.5D approaches including fine-tuning the models pre-trained from ImageNet as well as fine-tuning the 2D versions of our Models Genesis, confirming the importance of 3D anatomical information and significance of Models Genesis for 3D medical imaging. This performance is attributed to our unified self-supervised learning framework, built on a simple yet powerful observation: the sophisticated and recurrent anatomy in medical images can serve as strong yet free supervision signals for deep models to learn common anatomical representation automatically via self-supervision. As open science, all codes and pre-trained Models Genesis are available at <https://github.com/MrGiovanni/ModelsGenesis>.

© 2020 Elsevier B.V. All rights reserved.

1. Introduction

Transfer learning from *natural* images to *medical* images has become the *de facto* standard in deep learning for medical image analysis (Tajbakhsh et al., 2016; Shin et al., 2016), but given the marked differences between *natural* images and *medical* images, we hypothesize that transfer learning can yield more powerful (application-specific) *target* models from the *source* models built directly using medical images. To test this hypothesis, we have chosen chest imaging because the chest contains several critical organs, which are prone to a number of diseases that result in substantial morbidity and mortality, hence associated with significant health-care costs. In this research, we focus on Chest CT, because of its prominent role in diagnosing lung diseases, and our research community has accumulated several Chest CT image databases, for instance, LIDC-IDRI (Armato III et al., 2011) and NLST (NLST, 2011),

containing a large number of Chest CT images. However, systematically annotating Chest CT scans is not only tedious, laborious, and time-consuming, but it also demands costly, specialty-oriented skills, which are not easily accessible. Therefore, we seek to answer the following question: *Can we utilize the large number of available Chest CT images without systematic annotation to train source models that can yield high-performance target models via transfer learning?*

To answer this question, we have developed a framework that trains generic source models for 3D medical imaging. Our framework is *autodidactic*—eliminating the need for labeled data by self-supervision; *robust*—learning comprehensive image representation from a mixture of self-supervised tasks; *scalable*—consolidating a variety of self-supervised tasks into a single image restoration task with the same encoder-decoder architecture; and *generic*—benefiting a range of 3D medical imaging tasks through transfer learning. We call the models trained with our framework Generic Autodidactic Models, nicknamed Models Genesis, and refer to the model trained using Chest CT images as Genesis Chest CT. As ablation studies, we have also trained a downgraded 2D version using 2D Chest CT slices, called Genesis Chest CT 2D. For thorough per-

* Corresponding author.

E-mail addresses: Jianming.Liang@asu.edu, jianming.liang@computer.org (J. Liang).

Table 1

Pre-trained models with proxy tasks and target tasks. This paper uses transfer learning in a broader sense, where a *source model* is first trained to learn image presentation via *full supervision* or *self supervision* by solving a problem, called *proxy task* (general or application-specific), on a *source dataset* with *expert-provided* or *automatically-generated* labels, and then this *pre-trained* source model is fine tuned (transferred) through full supervision to yield a *target model* to solve application-specific problems (*target tasks*) in the same or different datasets (*target datasets*). We refer transfer learning to *same-domain* transfer learning when the models are pre-trained and fine-tuned within the same domain (modality, organ, disease, or dataset), and to *cross-domain* when the models are pre-trained in one domain and fine-tuned for a different domain.

Pre-trained model	Modality	Source dataset	Superv. / Annot.	Proxy task
Genesis Chest CT 2D	CT	LUNA 2016 (Setio et al., 2017)	Self / 0	Image restoration on 2D Chest CT slices
Genesis Chest CT (3D)	CT	LUNA 2016 (Setio et al., 2017)	Self / 0	Image restoration on 3D Chest CT volumes
Genesis Chest X-ray (2D)	X-ray	ChestX-ray8 (Wang et al., 2017b)	Self / 0	Image restoration on 2D Chest Radiographs
Models ImageNet	Natural	ImageNet (Deng et al., 2009)	Full / 14M images	Image classification on 2D ImageNet
Inflated 3D (I3D)	Natural	Kinetics (Carreira and Zisserman, 2017)	Full / 240K videos	Action recognition on human action videos
NiftyNet	CT	Pancreas-CT & BTCV (Gibson et al., 2018a)	Full / 90 cases	Organ segmentation on abdominal CT
MedicalNet	CT, MRI	3DSeg-8 (Chen et al., 2019b)	Full / 1638 cases	Disease/organ segmentation on 8 datasets
Code [†]	Object	Modality	Target dataset	Target task
NCC	Lung Nodule	CT	LUNA 2016 (Setio et al., 2017)	Lung nodule false positive reduction
NCS	Lung Nodule	CT	LIDC-IDRI (Armato III et al., 2011)	Lung nodule segmentation
ECC	Pulmonary Emboli	CT	PE-CAD (Tajbakhsh et al., 2015)	Pulmonary embolism false positive reduction
LCS	Liver	CT	LITS 2017 (Bilic et al., 2019)	Liver segmentation
BMS	Brain Tumor	MRI	BraTS 2018 (Menze et al., 2015; Bakas et al., 2018)	Brain tumor segmentation

[†] The first letter denotes the object of interest (“N” for lung nodule, “E” for pulmonary embolism, “L” for liver, etc); the second letter denotes the modality (“C” for CT, “M” for MRI, etc); the last letter denotes the task (“C” for classification, “S” for segmentation).

formance comparisons, we have trained a 2D model using Chest X-ray images, named as Genesis Chest X-ray (detailed in Table 1).

Naturally, 3D imaging tasks in the most prominent medical imaging modalities (e.g., CT and MRI) should be solved directly in 3D, but 3D models generally have significantly more parameters than their 2D counterparts, thus demanding more labeled data for training. As a result, learning from scratch simply in 3D may *not* necessarily yield performance better than fine-tuning Models ImageNet (i.e., pre-trained models on ImageNet), as revealed in Fig. 7. However, as demonstrated by our extensive experiments in Section 3, our Genesis Chest CT not only *significantly* outperforms learning 3D models from scratch (see Fig. 4), but also *consistently* tops any 2D/2.5D approaches including fine-tuning Models ImageNet as well as fine-tuning our Genesis Chest X-ray and Genesis Chest CT 2D (see Fig. 7 and Table 4). Furthermore, Genesis Chest CT surpasses publicly-available, pre-trained, (fully) supervised 3D models (see Table 3). Our results confirm the importance of 3D anatomical information and demonstrate the significance of Models Genesis for 3D medical imaging.

This performance is attributable to the following key observation: medical imaging protocols typically focus on particular parts of the body for specific clinical purposes, resulting in images of similar anatomy. The sophisticated yet recurrent anatomy offers consistent patterns for self-supervised learning to discover common representation of a particular body part (the lungs in our case). As illustrated in Fig. 1, the fundamental idea behind our self-supervised learning method is to recover anatomical patterns from images transformed via various ways in a unified framework.

In summary, we make the following three contributions:

1. A collection of generic pre-trained 3D models, performing effectively across diseases, organs, and modalities.
2. A scalable self-supervised learning framework, offering encoder for classification and encoder-decoder for segmentation.
3. A set of self-supervised training schemes, learning robust representation from multiple perspectives.

In the remainder of this paper, we first in Section 2 introduce our self-supervised learning framework for training Models Genesis, covering our four proposed image transformations with their learning perspectives, and describing the four unique properties of our Models Genesis. Section 3 details the training process of Mod-

els Genesis and the five target tasks for evaluating Models Genesis, while Section 4 summarizes the five major observations from our extensive experiments, demonstrating that our Models Genesis can serve as a primary source of transfer learning for 3D medical imaging. In Section 5, we discuss various aspects of Models Genesis, including their relationship with automated data augmentation, their impact on the creation of a medical ImageNet, and their capabilities for same- and cross-domain transfer learning followed by a thorough review of existing supervised and self-supervised representation learning approaches in medical imaging in Section 6. Finally, Section 7 concludes and outlines future extensions of Models Genesis.

2. Models Genesis

The objective of Models Genesis is to learn a common image representation that is transferable and generalizable across diseases, organs, and modalities. Fig. 1 depicts our self-supervised learning framework, which enables training 3D models from scratch using unlabeled images, consisting of three steps: (1) cropping sub-volumes from patient CT images, (2) deforming the sub-volumes, and (3) training a model to restore the original sub-volume. In the following sections, we first introduce the denotations of our self-supervised learning framework and then detail each of the training schemes with its learning objectives and perspectives, followed by a summary of the four unique properties of our Models Genesis.

2.1. Models Genesis learn by image restoration

Given a raw dataset consisting of N patient volumes, theoretically we can crop infinite number of sub-volumes from the dataset. In practice, we randomly generate a subset $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$, which includes n number of sub-volumes and then apply image transformation function to these sub-volumes, yielding

$$\tilde{\mathcal{X}} = f(\mathcal{X}), \quad (1)$$

where $\tilde{\mathcal{X}} = \{\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2, \dots, \tilde{\mathbf{x}}_n\}$ and $f(\cdot)$ denotes a transformation function. Subsequently, a Model Genesis, being an encoder-decoder network with skip connections in between, will learn to approx-

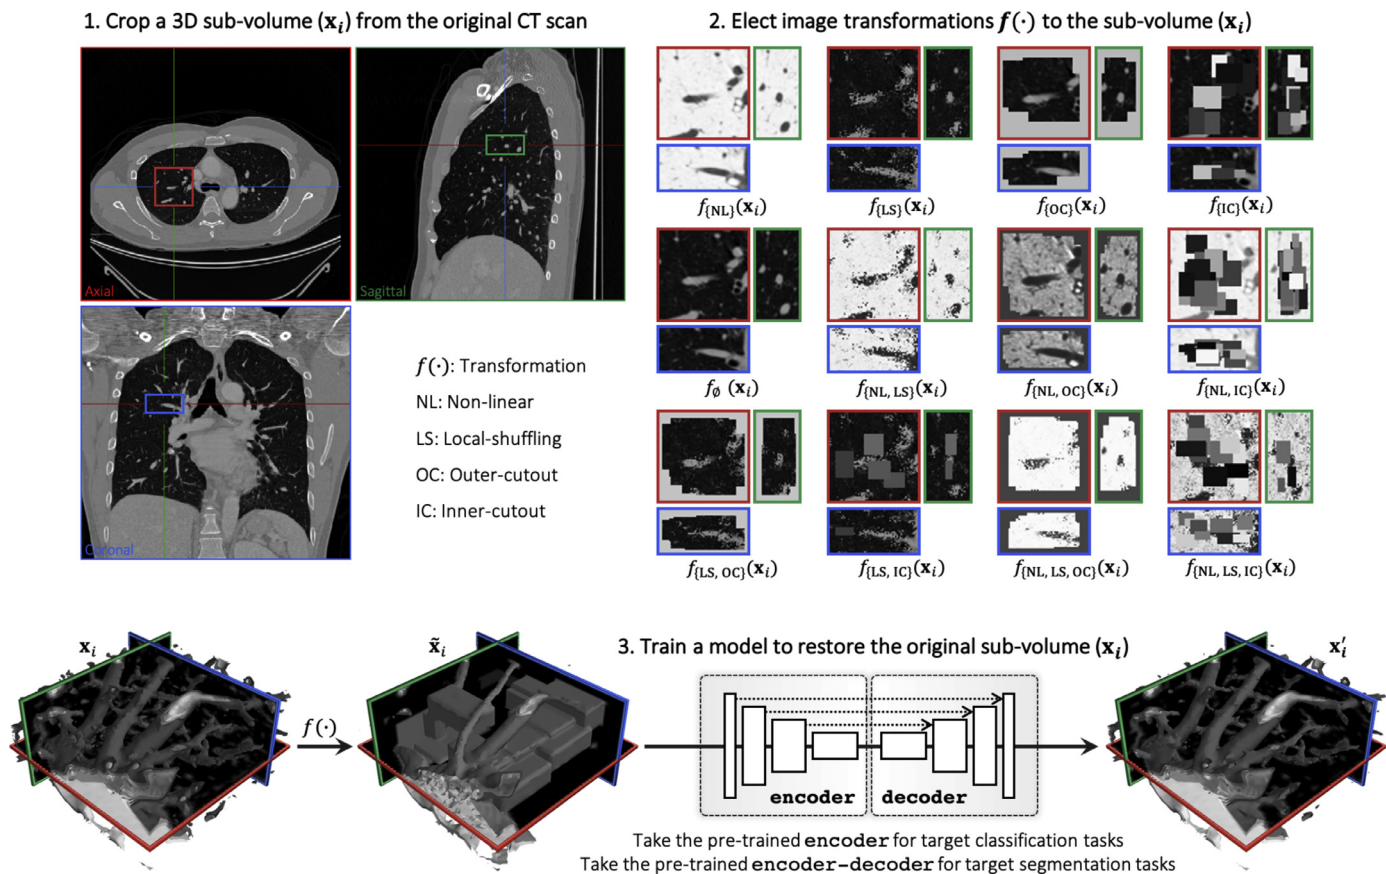


Fig. 1. [Better viewed on-line, in color, and zoomed in for details] Our self-supervised learning framework aims to learn general-purpose image representation by recovering the original sub-volumes of images from their transformed ones. We first crop arbitrarily-size sub-volume x_i at a random location from an unlabeled CT image. Each sub-volume x_i can undergo at most three out of four transformations: non-linear, local-shuffling, outer-cutout, and inner-cutout, resulting in a transformed sub-volume \tilde{x}_i . It should be noted that outer-cutout and inner-cutout are considered mutually exclusive. Therefore, in addition to the four original individual transformations, this process yields eight more transformations, including one identity mapping (ϕ meaning none of the four individual transformations is selected) and seven combined transformations. A Model Genesis, an encoder-decoder architecture with skip connections in between, is trained to learn a common image representation by restoring the original sub-volume x_i (as ground truth) from the transformed one \tilde{x}_i (as input), in which the reconstruction loss (MSE) is computed between the model prediction x'_i and ground truth x_i . Once trained, the encoder alone can be fine-tuned for target classification tasks; while the encoder and decoder together can be fine-tuned for target segmentation tasks.

imate the function $g(\cdot)$ which aims to map the transformed sub-volumes $\tilde{\mathcal{X}}$ back to their original ones \mathcal{X} , that is,

$$g(\tilde{\mathcal{X}}) = \mathcal{X} = f^{-1}(\tilde{\mathcal{X}}). \tag{2}$$

To avoid heavy weight dedicated towers for each proxy task and to maximize parameter sharing in Models Genesis, we consolidate four self-supervised schemes into a single image restoration task, enabling models to learn robust image representation by restoring from various sets of image transformations. Our proposed framework includes four transformations: (1) non-linear, (2) local-shuffling, (3) outer-cutout, and (4) inner-cutout. Each transformation is independently applied to a sub-volume with a predefined probability, while outer-cutout and inner-cutout are considered mutually exclusive. Consequently, each sub-volume can undergo at most three of the above transformations, resulting in twelve possible transformed sub-volume (see step 2 in Fig. 1). For clarity, we further define a *training scheme* as the process that (1) transforms sub-volumes using any of the aforementioned transformations, and (2) trains a model to restore the original sub-volumes from the transformed ones. For convenience, we refer to an *individual training scheme* as the scheme using one particular individual transformation. We should emphasize that our ultimate goal is not the task of image restoration *per se*. While restoring images is advocated and investigated as a training scheme for models to learn

image representation, the usefulness of the learned representation must be assessed *objectively* based on its generalizability and transferability to various target tasks.

2.2. Models Genesis learn from multiple perspectives

1) Learning appearance via non-linear transformation. We propose a novel self-supervised training scheme based on non-linear translation, with which the model learns to restore the intensity values of an input image transformed with a set of non-linear functions. The rationale is that the absolute intensity values (*i.e.*, Hounsfield units) in CT scans or relative intensity values in other imaging modalities convey important information about the underlying structures and organs (Buzug, 2011; Forbes, 2012). Hence, this training scheme enables the model to learn the appearance of the anatomic structures present in the images. In order to keep the appearance of the anatomic structures perceivable, we intentionally retain the non-linear intensity transformation function as *monotonic*, allowing pixels of different values to be assigned with new distinct values. To realize this idea, we use Bézier Curve (Mortenson, 1999), a smooth and monotonic transformation function, which is generated from two end points (P_0 and P_3) and two

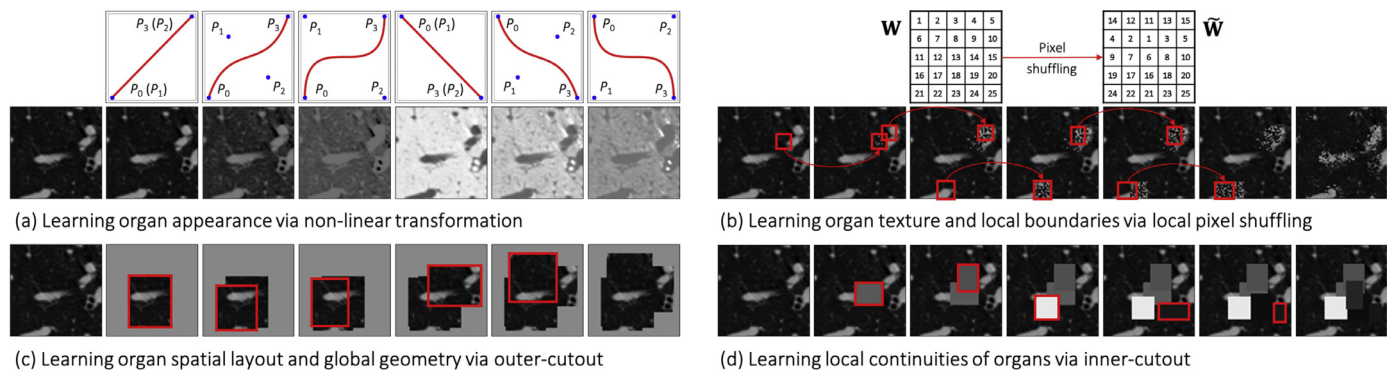


Fig. 2. Illustration of the proposed image transformations and their learning perspectives. For simplicity and clarity, we illustrate the transformation on a 2D CT slice, but our Genesis Chest CT is trained directly using 3D sub-volumes, which are transformed in a 3D manner. For ease of understanding, in (a) non-linear transformation, we have displayed an image undergoing different translating functions in Columns 2–7; in (b) local-shuffling, (c) outer-cutout, and (d) inner-cutout transformation, we have illustrated each of the processes step by step in Columns 2–6, where the first and last columns denote the original images and the final transformed images, respectively. In local-shuffling, a different window \mathbf{W} is automatically generated and used in each step. We provide the implementation details in Section 2.2 and more visualizations in Fig. D.13.

control points (P_1 and P_2), defined as:

$$B(t) = (1-t)^3 P_0 + 3(1-t)^2 t P_1 + 3(1-t)t^2 P_2 + t^3 P_3, \quad t \in [0, 1], \quad (3)$$

where t is a fractional value along the length of the line. In Fig. 2(a), we illustrate the original CT sub-volume (the left-most column) and its transformed ones based on different transformation functions. The corresponding transformation functions are shown in the top row. Notice that, when $P_0 = P_1$ and $P_2 = P_3$ the Bézier Curve is a linear function (shown in Columns 2, 5). Besides, we set $P_0 = (0, 0)$ and $P_3 = (1, 1)$ to get an increasing function (shown in Columns 2–4) and the opposite to get a decreasing function (shown in Columns 5–7). The control points are randomly generated for more variances (shown in Columns 3, 4, 6, 7). Before applying the transformation functions, in Genesis CT, we first clip the Hounsfield units values within the range of $[-1000, 1000]$ and then normalize each CT scan to $[0, 1]$, while in Genesis X-ray, we directly normalize each X-ray to $[0, 1]$ without intensity clipping.

2) Learning texture via local pixel shuffling. We propose local pixel shuffling to enrich local variations of a sub-volume without dramatically compromising its global structures, which encourages the model to learn the *local* boundaries and textures of objects. To be specific, for each input sub-volume, we randomly select 1000 windows and then shuffle the pixels inside each window sequentially. Mathematically, let us consider a small window \mathbf{W} with a size of $m \times n$. The local-shuffling acts on each window and can be formulated as

$$\tilde{\mathbf{W}} = \mathbf{P} \times \mathbf{W} \times \mathbf{P}', \quad (4)$$

where $\tilde{\mathbf{W}}$ is the transformed window, \mathbf{P} and \mathbf{P}' denote permutation metrics with the size of $m \times m$ and $n \times n$, respectively. Pre-multiplying \mathbf{W} with \mathbf{P} permutes the rows of the window \mathbf{W} , whereas post-multiplying \mathbf{W} with \mathbf{P}' results in the permutation of the columns of the window \mathbf{W} . The size of the local window determines the difficulty of proxy task. In practice, to preserve the global content of the image, we keep the window sizes smaller than the receptive field of the network, so that the network can learn much more robust image representation by “resetting” the original pixels positions. Note that our method is quite different from PatchShuffling (Kang et al., 2017), which is a regularization technique to avoid over-fitting. Unlike de-noising (Vincent et al., 2010) and in-painting (Pathak et al., 2016; Iizuka et al., 2017),

our local-shuffling transformation does not intend to replace the pixel values with noise, which therefore preserves the identical global distributions to the original sub-volume. In addition, local-shuffling within an extent keeps the objects perceivable, as shown in Fig. 2(b), benefiting the deep neural network in learning *local* invariant image representations, which serves as a complementary perspective with global patch shuffling (Chen et al., 2019a) (discussed in-depth in C1).

3) Learning context via outer and inner cutouts. We devise outer-cutout as a new training scheme for self-supervised learning. To realize it, we generate an arbitrary number (≤ 10) of windows, with various sizes and aspect ratios, and superimpose them on top of each other, resulting in a single window of a complex shape. When applying this merged window to a sub-volume, we leave the sub-volume region inside the window exposed and mask its surrounding (*i.e.*, outer-cutout) with a random number. Moreover, to prevent the task from being too difficult or even unsolvable, we extensively search for the optimal size of cutout regions spanning from 0% to 90%, incremented by 10% (detailed study presented in Appendix C3). In the end, we limit the outer-cutout region to be less than 1/4 of the whole sub-volume. By restoring the outer-cutouts, the model will learn the *global* geometry and spatial layout of organs in medical images via extrapolating within each sub-volume. We have illustrated this process step by step in Fig. 2(c). The first and last columns denote the original sub-volumes and the final transformed sub-volumes, respectively.

Our self-supervised learning framework also utilizes inner-cutout as a training scheme, where we mask the inner window regions (*i.e.*, inner-cutouts) and leave their surroundings exposed. By restoring the inner-cutouts, the model will learn *local* continuities of organs in medical images via interpolating within each sub-volume. Unlike Pathak et al. (2016), where in-painting is proposed as a proxy task by restoring only the central region of the image, we restore the entire sub-volume as the model output. Examples of inner-cutout are illustrated in Fig. 2(d). Following the suggestion from Pathak et al. (2016), the inner-cutout areas are limited to be less than 1/4 of the whole sub-volume, in order to keep the task reasonably difficult.

2.3. Models Genesis have several unique properties

1) Autodidactic—requiring no manual labeling. Models Genesis are trained in a self-supervised manner with abundant un-




Fig. 3. Comparing the combined training scheme with each of the proposed individual training schemes, we conduct statistical analyses between the top two training schemes as well as between the bottom two. Although some of the individual training schemes could be favorable for certain target tasks, there is no such clear clue to guarantee that any one of the individual training schemes would consistently offer the best performance on every target task. On the contrary, our combined training scheme consistently achieves the best results across all five target tasks.




Fig. 4. Models Genesis, as presented with the red vertical lines, achieve higher and more stable performance compared with three popular types of random initialization methods, including MSRA, Xavier, and Uniform. Among three out of the five applications, three different types of random distribution reveal no significant difference with respect to each other. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

labeled image datasets, demanding zero expert annotation effort. Consequently, Models Genesis are fundamentally different from traditional (fully) supervised transfer learning from ImageNet (Shin et al., 2016; Tajbakhsh et al., 2016), which offers modest benefits to 3D medical imaging applications as well as that from the existing pre-trained, full-supervised models including I3D (Carreira and Zisserman, 2017), NiftyNet (Gibson et al., 2018b), and MedicalNet (Chen et al., 2019b), which demand a volume of annotation effort to obtain the source models (statistics given in Table 1). To our best knowledge, this work represents the first effort to establish publicly-available, autodidactic models for 3D medical image analysis.

2) Robust—learning from multiple perspectives. Our combined approach trains Models Genesis from multiple perspectives (appearance, texture, context, etc.), leading to more robust models across all target tasks, as evidenced in Fig. 3, where our combined approach is compared with our individual schemes. This eclectic approach, incorporating multiple tasks into a single image restoration task, empowers Models Genesis to learn more comprehensive representation. While most self-supervised methods devise isolated training schemes to learn from specific perspectives—learning intensity value via colorization, context information via Jigsaw, orientation via rotation, etc—these methods are reported with mixed results on different tasks, in review papers such as Goyal et al. (2019), Kolesnikov et al. (2019), Taleb et al. (2020), and Jing and Tian (2020). It is critical as a multitude of state-of-the-art results in the literature show the importance of using compositions of more than one transformations per image (Graham, 2014; Dosovitskiy et al., 2015; Wu et al., 2020), which has also been experimentally confirmed in our image restoration task.

3) Scalable—accommodating many training schemes. Consolidated into a single image restoration task, our novel self-

supervised schemes share the same encoder and decoder during training. Had each task required its own decoder, due to limited memory on GPUs, our framework would have failed to accommodate a large number of self-supervised tasks. By unifying all tasks as a single image restoration task, any favorable transformation can be easily amended into our framework, overcoming the scalability issue associated with multi-task learning (Doersch and Zisserman, 2017; Noroozi et al., 2018; Standley et al., 2019; Chen et al., 2019b), where the network heads are subject to the specific proxy tasks.

4) Generic—yielding diverse applications. Models Genesis, trained via a diverse set of self-supervised schemes, learn a general-purpose image representation that can be leveraged for a wide range of target tasks. Specifically, Models Genesis can be utilized to initialize the encoder for the target classification tasks and to initialize the encoder-decoder for the target segmentation tasks, while the existing self-supervised approaches are largely focused on providing encoder models only (Jing and Tian, 2020). As shown in Table 3, Models Genesis can be generalized across diseases (e.g., nodule, embolism, tumor), organs (e.g., lung, liver, brain), and modalities (e.g., CT and MRI), a generic behavior that sets us apart from all previous works in the literature where the representation is learned via a specific self-supervised task, and thus lack generality.

3. Experiments

3.1. Pre-training Models Genesis

Our Models Genesis are pre-trained from 623 Chest CT scans in LUNA 2016 (Setio et al., 2017) in a self-supervised manner. The reason that we decided not to use all 888 scans provided by this dataset was to avoid test-image leaks between proxy and target tasks, so that we can confidently use the rest of the images solely

Table 2

Genesis CT is pre-trained on only LUNA 2016 dataset (*i.e.*, the source) and then fine-tuned for five distinct medical image applications (*i.e.*, the targets). These target tasks are selected such that they show varying levels of semantic distance from the source, in terms of organs, diseases, and modalities, allowing us to investigate the transferability of the pre-trained weights of Genesis CT with respect to the domain distance. The cells checked by \times denote the properties that are different between the source and target datasets.

Task	Disease	Organ	Dataset	Modality
NCC				
NCS				
ECC	\times		\times	
LCS	\times	\times	\times	
BMS	\times	\times	\times	\times

for testing Models Genesis as well as the target models, although Models Genesis are trained from only unlabeled images, involving no annotation shipped with the dataset. We first randomly crop sub-volumes, sized $64 \times 64 \times 32$ pixels, from different locations. To extract more informative sub-volumes for training, we then intentionally exclude those which are empty (air) or contain full tissues. Our Models Genesis 2D are self-supervised pre-trained from LUNA 2016 (Setio et al., 2017) and ChestX-ray14 (Wang et al., 2017b) using 2D CT slices in an axial view and X-ray images, respectively. For all proxy tasks and target tasks, the raw image intensities were normalized to the [0,1] range before training. We use the mean square error (MSE) between input and output images as objective function for the proxy task of image restoration. As suggested by Pathak et al. (2016) and Chen et al. (2019a), the MSE loss is sufficient for representation learning, although the restored images may be blurry.

When pre-training Models Genesis, we apply each of the transformations on sub-volumes with a pre-defined probability. That

being said, the model will encounter not only the transformed sub-volumes as input, but also the original sub-volumes. This design offers two advantages:

- First, the model must distinguish original versus transformed images, discriminate transformation type(s), and restore images if transformed. Our self-supervised learning framework, therefore, results in pre-trained models that are capable of handling versatile tasks.
- Second, since original images are presented in the proxy task, the semantic difference of input images between the proxy and target task becomes smaller. As a result, the pre-trained model can be transferable to process regular/normal images in a broad variety of target tasks.

3.2. Fine-tuning Models Genesis

The pre-trained Models Genesis can be adapted to new imaging tasks through transfer learning or fine-tuning. There are three major transfer learning scenarios: (1) employing the encoder as a fixed feature extractor for a new dataset and following up with a linear classifier (*e.g.*, Linear SVM or Softmax classifier), (2) taking the pre-trained encoder and appending a sequence of fully-connected (*fc*) layers for target classification tasks, and (3) taking the pre-trained encoder and decoder and replacing the last layer with a $1 \times 1 \times 1$ convolutional layer for target segmentation tasks. For scenarios (2) and (3), it is possible to fine-tune all the layers of the model or to keep some of the earlier layers fixed, only fine-tuning some higher-level portion of the model. We have evaluated the performance of our self-supervised representation for transfer learning by fine-tuning all layers in the network. In the following, we examine Models Genesis on five distinct medical applications, covering classification and segmentation tasks in CT and MRI images with varying levels of semantic distance from the source (Chest CT) to the targets in terms of organs, diseases, and modal-

Table 3

Our Models Genesis lead the best or comparable performance on five distinct medical target tasks over six self-supervised learning approaches (revised in 3D) and three competing publicly available (fully) supervised pre-trained 3D models. For ease of comparison, we evaluate AUC score for the two classification tasks (*i.e.*, NCC and ECC) and IoU score for the three segmentation tasks (*i.e.*, NCS, LCS, and BMS). All of the results, including the mean and standard deviation (mean \pm s.d.) across ten trials, reported in the table are evaluated using our dataset splitting, elaborated in Section 3.2. For every target task, we have further performed independent two sample *t*-test between the best (bolded) vs. others and highlighted boxes in blue when they are not statistically significantly different at $p = .05$ level. The footnotes compare our results with the state-of-the-art performance for each target task, using the evaluation metric for the data acquired from competitions.

Pre-training	Approach	Target tasks				
		NCC ^a (%)	NCS ^b (%)	ECC ^c (%)	LCS ^d (%)	BMS ^e (%)
No	Random with Uniform Init	94.74 \pm 1.97	75.48 \pm 0.43	80.36 \pm 3.58	78.68 \pm 4.23	60.79 \pm 1.60
	Random with Xavier Init (Glorot and Bengio, 2010)	94.25 \pm 5.07	74.05 \pm 1.97	79.99 \pm 8.06	77.82 \pm 3.87	58.52 \pm 2.61
	Random with MSRA Init (He et al., 2015)	96.03 \pm 1.82	76.44 \pm 0.45	78.24 \pm 3.60	79.76 \pm 5.43	63.00 \pm 1.73
(Fully) supervised	I3D (Carreira and Zisserman, 2017)	98.26\pm0.27	71.58 \pm 0.55	80.55 \pm 1.11	70.65 \pm 4.26	67.83\pm0.75
	NiftyNet (Gibson et al., 2018b)	94.14 \pm 4.57	52.98 \pm 2.05	77.33 \pm 8.05	83.23 \pm 1.05	60.78 \pm 1.60
	MedicalNet (Chen et al., 2019b)	95.80 \pm 0.49	75.68 \pm 0.32	86.43\pm1.44	85.52\pm0.58^f	66.09 \pm 1.35
Self-supervised	De-noising (Vincent et al., 2010)	95.92 \pm 1.83	73.99 \pm 0.62	85.14\pm3.02	84.36 \pm 0.96	57.83 \pm 1.57
	In-painting (Pathak et al., 2016)	91.46 \pm 2.97	76.02 \pm 0.55	79.79 \pm 3.55	81.36 \pm 4.83	61.38 \pm 3.84
	Jigsaw (Noroozi and Favaro, 2016)	95.47 \pm 1.24	70.90 \pm 1.55	81.79 \pm 1.04	82.04 \pm 1.26	63.33 \pm 1.11
	DeepCluster (Caron et al., 2018)	97.22 \pm 0.55	74.95 \pm 0.46	84.82 \pm 0.62	82.66 \pm 1.00	65.96 \pm 0.85
	Patch shuffling (Chen et al., 2019a)	91.93 \pm 2.32	75.74 \pm 0.51	82.15 \pm 3.30	82.82 \pm 2.35	52.95 \pm 6.92
	Rubiks Cube (Zhuang et al., 2019)	96.24 \pm 1.27	72.87 \pm 0.16	80.49 \pm 4.64	75.59 \pm 0.20	62.75 \pm 1.93
	Genesis Chest CT (ours)	98.34\pm0.44	77.62\pm0.64	87.20\pm2.87	85.10\pm2.15	67.96\pm1.29

^a The winner in LUNA (2016) holds an official score of 0.968 vs. 0.971 (ours).

^b Wu et al. (2018) holds a Dice of 74.05% vs. 75.86% \pm 0.90% (ours).

^c Zhou et al. (2017) holds an AUC of 87.06% vs. 87.20% \pm 2.87% (ours).

^d The winner in LiTS (2017) with post-processing holds a Dice of 96.60% vs. 93.19% \pm 0.46% (ours without post-processing).

^e MRI Flair images are only utilized for segmenting brain tumors, so the results are not submitted to BraTS 2018.

^f Genesis Chest CT is slightly outperformed by MedicalNet in LCS because the latter has been (fully) supervised pre-trained on the LiTS dataset.

Table 4

Our 3D approach, initialized by Models Genesis, significantly elevates the classification performance compared with 2.5D and 2D approaches in reducing lung nodule and pulmonary embolism false positives. The entries in bold highlight the best results achieved by different approaches. For the 2D slice-based approach, we extract input consisting of three adjacent axial views of the lung nodule or pulmonary embolism and some of their surroundings. For the 2.5D orthogonal approach, each input is composed of an axial, coronal, and sagittal slice and centered at a lung nodule or pulmonary embolism candidate.

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

ities (see Table 2) for investigating the transferability of Models Genesis.

3.2.1. Lung nodule false positive reduction (NCC)

The dataset is provided by LUNA 2016 (Setio et al., 2017) and consists of 888 low-dose lung CTs with slice thickness less than 2.5mm. Patients are randomly assigned into a training set (445 cases), a validation set (178 cases), and a test set (265 cases). The dataset offers the annotations for a set of 5,510,166 candidate locations for the false positive reduction task, wherein true positives are labeled as “1” and false positives are labeled as “0”. Following the prior works (Setio et al., 2016; Sun et al., 2017c), we evaluate performance via Area Under the Curve (AUC) score on classifying true positives and false positives.

3.2.2. Lung nodule segmentation (NCS)

The dataset is provided by the Lung Image Database Consortium image collection (LIDC-IDRI) (Armato III et al., 2011) and consists of 1018 cases collected by seven academic centers and eight medical imaging companies. The cases were split into training (510), validation (100), and test (408) sets. Each case is a 3D CT scan and the nodules have been marked as volumetric binary masks. We have re-sampled the volumes to 1-1-1 spacing and then extracted a $64 \times 64 \times 32$ crop around each nodule. These 3D crops are used for model training and evaluation. As in prior works (Aresta et al., 2019; Tang et al., 2019; Zhou et al., 2018), we adopt Intersection over Union (IoU) and Dice coefficient scores to evaluate performance. Note that for this particular application, we calculate mean of the IoUs at thresholds ranging from 0.5 to 0.95 with a step size of 0.05.

3.2.3. Pulmonary embolism false positive reduction (ECC)

We utilize a database consisting of 121 computed tomography pulmonary angiography (CTPA) scans with a total of 326 emboli. Following the prior works (Liang and Bi, 2007), we utilize their PE candidate generator based on the toboggan algorithm, resulting in total of 687 true positives and 5568 false positives. The dataset is then divided at the patient-level into a training set with 434 true positive PE candidates and 3406 false positive PE candidates, and a test set with 253 true positive PE candidates and 2162 false positive PE candidates. To conduct a fair comparison with the prior study (Zhou et al., 2017; Tajbakhsh et al., 2016; 2019b), we compute candidate-level AUC on classifying true positives and false positives.

3.2.4. Liver segmentation (LCS)

The dataset is provided by MICCAI 2017 LiTS Challenge and consists of 130 labeled CT scans, which we split into training (100

patients), validation (15 patients), and test (15 patients) subsets. The ground truth segmentation provides two different labels: liver and lesion. For our experiments, we only consider liver as positive class and others as negative class and evaluate segmentation performance using Intersection over Union (IoU) and Dice coefficient scores.

3.2.5. Brain tumor segmentation (BMS)

The dataset is provided by BraTS 2018 challenge (Menze et al., 2015; Bakas et al., 2018) and consists of 285 patients (210 HGG and 75 LGG), each with four 3D MRI modalities (T1, T1c, T2, and Flair) rigidly aligned. We adopt 3-fold cross validation, in which two folds (190 patients) are for training and one fold (95 patients) for test. Annotations include background (label 0) and three tumor subregions: GD-enhancing tumor (label 4), the peritumoral edema (label 2), and the necrotic and non-enhancing tumor core (label 1). We consider those with label 0 as negatives and others as positives and evaluate segmentation performance using Intersection over Union (IoU) and Dice coefficient scores.

3.3. Benchmarking Models Genesis

For a thorough comparison, we used three different techniques to randomly initialize the weights of models: (1) a basic random initialization method based on Gaussian distributions, (2) a method commonly known as Xavier, which was suggested in Glorot and Bengio (2010), and (3) a revised version of Xavier called MSRA, which was suggested in He et al. (2015). They are implemented as `uniform`, `glorot_uniform`, and `he_uniform`, respectively, following the Initializers¹ in Keras. We compare Models Genesis with Rubik’s cube (Zhuang et al., 2019), the most recent multi-task and self-supervised learning method for 3D medical imaging. Considering that most of the self-supervised learning methods are initially proposed and implemented in 2D, we have extended five most representative ones (Vincent et al., 2010; Pathak et al., 2016; Noroozi and Favaro, 2016; Chen et al., 2019a; Caron et al., 2018) into their 3D versions for a fair comparison (see detailed implementation in Appendix A). To promote the 3D self-supervised learning research, we make our own implementation of the 3D extended methods and their corresponding pre-trained weights publicly available as an open-source tool that can effectively be used out-of-the-box. In addition, we have examined publicly available pre-trained models for 3D transfer learning in medical imaging, including NiftyNet² (Gibson et al., 2018b), Medical-

¹ Initializers: [faroit.com/keras-docs/1.2.2/initializers](https://keras.io/initializers/)

² NiftyNet Model Zoo: github.com/NifTK/NiftyNetModelZoo

Net³ (Chen et al., 2019b), and, the most influential 2D weights initialization, Models ImageNet. We also fine-tune I3D⁴ (Carreira and Zisserman, 2017) in our five target tasks because it has been shown to successfully initialize 3D models for lung nodule detection in Ardila et al. (2019). The detailed configurations of these models can be found in Appendix B.

3D U-Net architecture⁵ is used in 3D applications; U-Net architecture⁶ is used in 2D applications. Batch normalization (Ioffe and Szegedy, 2015) is utilized in all 3D/2D deep models. For proxy tasks, SGD method (Zhang, 2004) with an initial learning rate of $1e-5$ is used for optimization. We use ReduceLROnPlateau to schedule learning rate, in which if no improvement is seen in the validation set for a certain number of epochs, the learning rate is reduced. For target tasks, Adam method (Kinga and Adam, 2015) with a learning rate of $1e-3$ is used for optimization, where $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-8$. We use *early-stop* mechanism on the validation set to avoid over-fitting. Simple yet heavy 3D data augmentation techniques are employed in all five target tasks, including random flipping, transposing, rotating, and adding Gaussian noise. We run each method ten times on all of the target tasks and report the average, standard deviation, and further present statistical analysis based on an independent two-sample *t*-test.

In the proxy task, we pre-train the model using 3D sub-volumes sized $64 \times 64 \times 32$, whereas in target tasks, the input is not limited to sub-volumes with certain size. That being said, our pre-trained models can be fine-tuned in the tasks with CT sub-volumes, entire CT volumes, or even MRI volumes as input upon user's need. The flexibility of input size is attributed to two reasons: (1) our pre-trained models learn generic image representation such as appearance, texture, and context feature, and (2) the encoder-decoder architecture is able to process images with arbitrary sizes.

4. Results

In this section, we begin with an ablation study to compare the combined approach with each individual scheme, concluding that the combined approach tends to achieve more robust results and consistently exceeds any other training schemes. We then take our pre-trained model from the combined approach and present results on five 3D medical applications, comparing them against the state-of-the-art approaches found in recent supervised and self-supervised learning literature.

4.1. The combined learning scheme exceeds each individual

We have devised four individual training schemes by applying each of the transformations (*i.e.*, non-linear, local-shuffling, outer-cutout, and inner-cutout) individually to a sub-volume and training the model to restore the original one. We compare each of these training schemes with identical-mapping, which does not involve any image transformation. In three out of the five target tasks, as shown in Figs. 3–4, the model pre-trained by identical-mapping scheme does not perform as well as random initialization. This undesired representation obtained via identical-mapping suggests that without any image transformation, the model would not benefit much from the proxy image restoration task. On the contrary, nearly all of the individual schemes offer higher target task performances than identical-mapping, demonstrating the significance of the four devised image transformations in learning image representation.

Although each of the individual schemes has established the capability in learning image representation, its empirical performance varies from task to task. That being said, given a target task, there is no clear winner among the four individual schemes that can always guarantee the highest performance. As a result, we have further devised a combined scheme, which applies transformations to a sub-volume with a predefined probability for each transformation and trains a model to restore the original one. To demonstrate the importance of combining these image transformations together, we examine the combined training scheme against each of the individual ones. Fig. 3 shows that the combined scheme consistently exceeds any other individual schemes in all five target tasks. We have found that the combination of different transformations is advantageous because, as discussed, we cannot rely on one single training scheme to achieve the most robust and compelling results across multiple target tasks. It is our novel representation learning framework based on image restoration that allows integrating various training schemes into a single training scheme. Our qualitative assessment of image restoration quality, provided in Fig. D.14, further indicates that the combined scheme is superior over all four individual schemes in restoring the images that have been undergone multiple transformations. In summary, our combined scheme pre-trains a model from multiple perspectives (appearance, texture, context, etc.), empowering models to learn a more comprehensive representation, thereby leading to more robust target models. Based on the above ablation studies, in the following sections, we refer the models pre-trained by the combined scheme to Models Genesis and, in particular, refer the model pre-trained on LUNA 2016 dataset to Genesis Chest CT.

4.2. Models Genesis outperform learning from scratch

Transfer learning accelerates training and boosts performance, only if the image representation learned from the original (proxy) task is general and transferable to target tasks. Fine-tuning models trained on ImageNet has been a great success story in 2D (Tajbakhsh et al., 2016; Shin et al., 2016), but for 3D representation learning, there is no such a massive labeled dataset like ImageNet. As a result, it is still common practice to train 3D model from scratch in 3D medical imaging. Therefore, to establish the 3D baselines, we have trained 3D models with three representative random initialization methods, including naive uniform initialization, Xavier/Glorot initialization proposed by Glorot and Bengio (2010), and He normal (MSRA) initialization proposed by He et al. (2015). When comparing deep model initialization by transfer learning and by controlling mathematical distribution, the former learns more sophisticated image representation but suffers from a domain gap, whereas the latter is task independent yet provides relatively less benefit than the former. The hypothesis underneath transfer learning is that transferring deep features across visual tasks can obtain a semantically more powerful representation, compared with simply initializing weights using different distributions. From our comprehensive experiments in Fig. 4, we have observed the following:

- Within each method, random initialization of weights has shown large variance in results of ten trials; it is in large part due to the difficulty of adequately initializing these networks from scratch. A small miscalibration of the initial weights can lead to vanishing or exploding gradients, as well as poor convergence properties.
- In three out of the five 3D medical applications, the results reveal no significant difference among these random initialization methods. Although randomly initializing weights can vary by the behaviors on different applications, He normal (MSRA), in which the weights are initialized with a specific ReLU-aware

³ MedicalNet: github.com/Tencent/MedicalNet

⁴ I3D: github.com/deepmind/kinetics-i3d

⁵ 3D U-Net: github.com/ellisdg/3DUNetCNN

⁶ Segmentation Models: github.com/qubvel/segmentation_models




Fig. 5. Models Genesis enable better optimization than learning from scratch, evident by the learning curves for the target tasks of reducing false positives in detecting lung nodules (NCC) and pulmonary embolism (ECC) as well as segmenting lung nodule (NCS), liver (LCS), and brain tumor (BMS). We have plotted the validation performance averaged by ten trials for each application, in which accuracy and dice-coefficient scores are reported for classification and segmentation tasks, respectively. As seen, initializing with our pre-trained Models Genesis demonstrates benefits in the convergence speed.

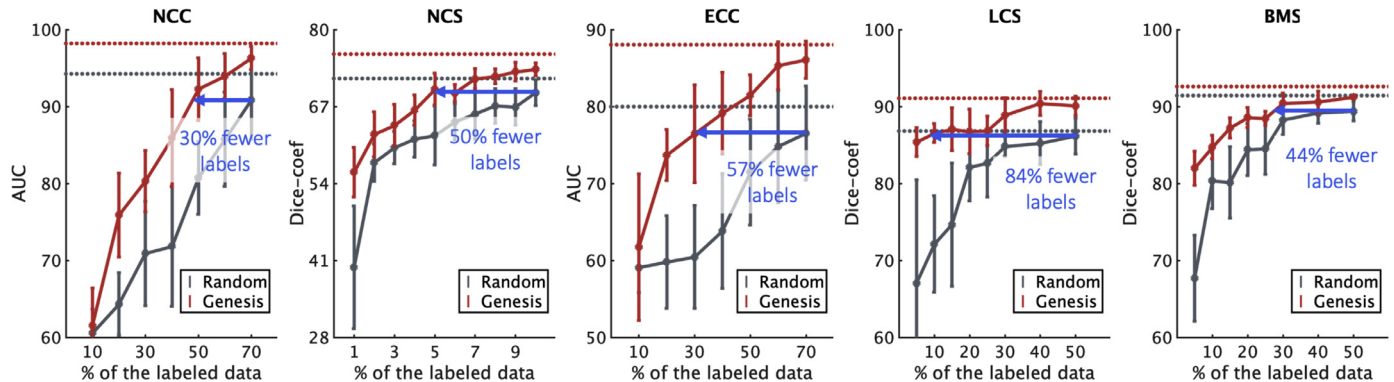


Fig. 6. Initializing with our Models Genesis, the annotation cost can be reduced by 30%, 50%, 57%, 84%, and 44% for target tasks NCC, NCS, ECC, LCS, and BMS, respectively. With decreasing amounts of labeled data, Models Genesis (red) retain a much higher performance on all five target tasks, whereas learning from scratch (grey) fails to generalize. Note that the horizontal red and gray lines refer to the performances that can eventually be achieved by Models Genesis and learning from scratch, respectively, when using the entire dataset. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

initialization, generally works the most reliably among all five target tasks.

- On the other hand, initialization with our pre-trained Genesis Chest CT stabilizes the overall performance and, more importantly, elevates the average performance over all three random initialization methods by a large margin. Our statistical analysis shows that the performance gain is significant for all the target tasks under study. This suggests that, owing to the representation learning scheme, our initial weights provide a better starting point than the ones generated under particular statistical distributions, while being over 13% faster (see Fig. 5). This observation has also been widely obtained in 2D model initialization (Tajbakhsh et al., 2016; Shin et al., 2016; Rawat and Wang, 2017; Zhou et al., 2017; Voulodimos et al., 2018).

Altogether, in contrast to 3D scratch models, we believe Models Genesis can potentially serve as a primary source of transfer learning for 3D medical imaging applications. Besides contrasting with the three random initialization methods, we further examine our Models Genesis against the existing pre-trained 3D models in the coming section.

4.3. Models Genesis surpass existing pre-trained 3D models

We have evaluated our Models Genesis with existing publicly available pre-trained 3D models on five distinct medical target tasks. As shown in Table 3, Genesis Chest CT noticeably contrasts with any other existing 3D models, which have been pre-trained by full supervision. Note that, in the liver segmentation task (LCS), Genesis Chest CT is slightly outperformed by Medical-

Net because of the benefit that MedicalNet gained from its (fully) supervised pre-training on the LiTS dataset directly. Further statistical tests reveal that Genesis Chest CT still yields comparable performance with MedicalNet at $p = .05$ level. For the rest four target tasks, Genesis Chest CT achieves superior performance against all its counterparts by a large margin, demonstrating the effectiveness and transferability of the learned features of Models Genesis, which are beneficial for both classification and segmentation tasks.

More importantly, although Genesis Chest CT is pre-trained on Chest CT only, it can generalize to different organs, diseases, datasets, and even modalities. For instance, the target task of pulmonary embolism false positive reduction is performed in Contrast-Enhanced CT scans that can appear differently from the proxy tasks in normal CT scans; yet, Genesis Chest CT achieves a remarkable improvement over training from scratch, increasing the AUC by 7 points. Moreover, Genesis Chest CT continues to yield a significant IoU gain in liver segmentation even though the proxy task and target task are significantly different in both, diseases affecting the organs (lung vs. liver) and the dataset itself (LUNA 2016 vs. LiTS 2017). We have further examined Genesis Chest CT and other existing pre-trained models using MRI Flair images, which represent the widest domain distance between the proxy and target tasks. As reported in Table 3 (BMS), Genesis Chest CT yields nearly a 5-point improvement in comparison with random initialization. The increased performance on the MRI imaging task is a particularly strong demonstration of the transfer learning capabilities of our Genesis Chest CT. To further investigate the behavior of Genesis Chest CT when encountering medical images




Fig. 7. When solving problems in volumetric medical modalities, such as CT and MRI images, 3D *volume-based* approaches consistently offer superior performance than 2D *slice-based* approaches empowered by transfer learning. We conduct statistical analyses (circled in blue) between the highest performance achieved by 3D and 2D solutions. Training 3D models from scratch does not necessarily outperform their 2D counterparts (see NCC). However, training the same 3D models from Genesis Chest CT outperforms all their 2D counterparts, including fine-tuning Models ImageNet as well as fine-tuning our Genesis Chest X-ray and Genesis Chest CT 2D. It confirms the effectiveness of Genesis Chest CT in unlocking the power of 3D models. In addition, we have also provided statistical analyses between the highest and the second highest performances achieved by 2D models, finding that Models Genesis (2D) offer equivalent performances (n.s.) to Models ImageNet in four out of the five applications. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

from different modalities, we have provided extensive visualization in Fig. D.15, including example images from CT, X-ray, Ultrasound, and MRI modalities.

Considering the model footprint, our Models Genesis take the basic 3D U-Net as the backbone, carrying much fewer parameters than the existing open-source pre-trained 3D models. For example, we have adopted MedicalNet with resnet-101 as the backbone, which offers the highest performance based on Chen et al. (2019b) but comprises of 85.75M parameters; the pre-trained I3D (Carreira and Zisserman, 2017) contains 25.35M parameters in the encoder; the pre-trained NiftyNet uses Dense V-Networks (Gibson et al., 2018a) as backbone, comprising of only 2.60M parameters, but it does not perform as well as its counterparts in all five target tasks. Taken together, these results indicate that our Models Genesis, with only 16.32M parameters, surpass all existing pre-trained 3D models in terms of generalizability, transferability, and parameter efficiency.

4.4. Models Genesis reduce annotation efforts by at least 30%

While critics often stress the need for sufficiently large amounts of labeled data to train a deep model, transfer learning leverages the knowledge about medical images already learned by pre-trained models and therefore requires considerably fewer annotated data and training iterations than learning from scratch. We have simulated the scenarios of using a handful of labeled data, which allows investigating the power of our Models Genesis in transfer learning. Fig. 6 displays the results of training with a partial dataset, demonstrating that fine-tuning Models Genesis saturates quickly on the target tasks since it can achieve similar performance compared with the full dataset training. Specifically, the performance of learning 3D models from scratch with entire datasets can be approximated using Models Genesis with only 50%, 5%, 30%, 5%, and 30% of datasets for NCC, NCS, ECC, LCS, and BMS, respectively. This shows that our Models Genesis can mitigate

the lack of labeled images, resulting in a more annotation efficient deep learning in the end.

Furthermore, the performance gap between fine-tuning and learning from scratch is significant and steady over training models with each partial data point. For the lung nodule false positive reduction target task (NCC in Fig. 6), using only 49% training data, Models Genesis equal the performance of 70% training data learning from scratch. Therefore, about 30% of the annotation cost associated with learning from scratch in NCC is recovered by initializing with Models Genesis. For the lung nodule segmentation target task (NCS in Fig. 6), with 5% training data, Models Genesis can achieve the performance equivalent to learning from scratch using 10% training data. Based on this analysis, the cost of annotation in NCS can be reduced by half using Models Genesis compared with learning from scratch. For the pulmonary embolism false positive reduction target task (ECC), Fig. 6 suggests that with only 30% training samples, Models Genesis achieve performance equivalent to learning from scratch using 70% training samples. Therefore, nearly 57% of the labeling cost associated with the use of learning from scratch for ECC could be recovered with our Models Genesis. For the liver segmentation target task (LCS) in Fig. 6, using 8% training data, Models Genesis equal the performance of learning from scratch using 50% training samples. Therefore, about 84% of the annotation cost associated with learning from scratch in LCS is recovered by initializing with Models Genesis. For the brain tumor segmentation target task (BMS) in Fig. 6, with less than 28% training data, Models Genesis achieve the performance equivalent to learning from scratch using 50% training data. Therefore, nearly 44% annotation efforts can be reduced using Models Genesis compared with learning from scratch. Overall, at least 30% annotation efforts have been reduced by Models Genesis, in comparison with learning a 3D model from scratch in five target tasks. With such annotation-efficient 3D transfer learning paradigm, computer-aided diagnosis of rare diseases or rapid response to global pandemics, which are severely underrepresented owing to the difficulty of collecting a sizeable amount labeled data, could be eventually actualized.

4.5. Models Genesis consistently top any 2D/2.5D approaches

We have thus far presented the power of 3D models in processing volumetric data, in particular, with limited annotation. Besides adopting 3D models, another common strategy to handle limited data in volumetric medical imaging is to reformat 3D data into a 2D image representation followed by fine-tuning pre-trained Models ImageNet (Shin et al., 2016; Tajbakhsh et al., 2016). This approach increases the training examples by order of magnitude, but it sacrifices the 3D context. It is interesting to note how Genesis Chest CT compares with this *de facto* standard in 2D. We have thus implemented two different methods to reformat 3D data into 2D input: the regular 2D representation obtained by extracting adjacent axial slices (Ben-Cohen et al., 2016; Sun et al., 2017a), and the 2.5D representation (Prasoon et al., 2013; Roth et al., 2014; 2015) composed of axial, coronal, and sagittal slices from volumetric data. Both of these 2D approaches seek to use 2D representation to emulate something three dimensional, in order to fit the paradigm of fine-tuning Models ImageNet. In the inference, classification and segmentation tasks are evaluated differently in 2D: for classification, the model predicts labels of slices extracted from the center locations because other slices are not guaranteed to include objects; for segmentation, the model predicts segmentation mask slice by slice and form the 3D segmentation volume by simply stacking the 2D segmentation maps.

Fig. 7 exposes the comparison between 3D and 2D models on five 3D target tasks. Additionally, Table 4 compares 2D slice-based, 2.5D orthogonal, and 3D volume-based approaches on lung nodule and pulmonary embolism false positive reduction tasks. As evidenced by our statistical analyses, the 3D models trained from Genesis Chest CT achieve significantly higher average performance and lower standard deviation than 2D models fine-tuned from ImageNet using either 2D or 2.5D image representation. Nonetheless, the same conclusion does not apply to the models trained from scratch—3D scratch models are outperformed by 2D models in one out of the five target tasks (*i.e.*, NCC in Fig. 7 and Table 4) and also exhibit an undesirably larger standard deviation. We attribute the mixed results of 3D scratch models to the larger number of model parameters and limited sample size in the target tasks, which together impede the full utilization of 3D context. In fact, the undesirable performance of the 3D scratch models highlights the effectiveness of Genesis Chest CT, which unlocks the power of 3D models for medical imaging. To summarize, we believe that 3D problems in medical imaging should be solved in 3D directly.

5. Discussions

5.1. Do we still need a Medical ImageNet?

In computer vision, at the time this paper is written, no self-supervised learning method outperforms fine-tuning models pre-trained on ImageNet (Jing and Tian, 2020; Chen et al., 2019a; Kolesnikov et al., 2019; Zhou et al., 2019b; Hendrycks et al., 2019; Zhang et al., 2019; Caron et al., 2019). Therefore, it may seem surprising to observe from our results in Table 3 that (fully) supervised representation learning methods do not necessarily offer higher performances in some 3D target tasks than self-supervised representation learning methods. We ascribe this phenomenon to the limited amount of supervision used in their pre-training (90 cases for NiftyNet (Gibson et al., 2018b) and 1638 cases for MedicalNet (Chen et al., 2019b)) or the domain distance (from videos to CT/MRI for I3D (Carreira and Zisserman, 2017)). Evidenced by a prior study (Sun et al., 2017b) on ImageNet pre-training, large amount of supervision is required to foster a generic, comprehensive image representation. Back in 2009, when ImageNet had not been established, it was challenging to empower a deep model

with generic image representation using a small or even medium size of labeled data, the same situation, we believe, that presents in 3D medical image analysis today. Therefore, despite the outstanding performance of Models Genesis, there is no doubt that a large, strongly annotated dataset for medical image analysis, like ImageNet (Deng et al., 2009) for computer vision, is still highly demanded. One of our goals for developing Models Genesis is to help create such a medical ImageNet. Based on a small set of expert annotations, models fine-tuned from Models Genesis will be able to help quickly generate initial rough annotations of unlabeled images for expert review, thus reducing the annotation efforts and accelerating the creation of a large, strongly annotated, medical ImageNet. In summary, Models Genesis are not designed to replace such a large, strongly annotated dataset for medical image analysis, as ImageNet for computer vision, but rather to help create one.

5.2. Same-domain or cross-domain transfer learning?

Same-domain transfer learning is always preferred whenever possible because a relatively smaller domain gap makes the learned image representation more beneficial for target tasks. Even the most recent self-supervised learning approaches in medical imaging were solely evaluated within the same dataset, such as Chen et al. (2019a); Tajbakhsh et al. (2019a); Zhu et al. (2020). Same-domain transfer learning strikes as a preferred choice in terms of performance; however, most of the existing medical datasets, with less than hundred cases, are usually too small for deep models to learn reliable image representation. Therefore, for our future work, we plan to combine the publicly available datasets from similar domains together to train modality-oriented models, including Genesis CT, Genesis MRI, Genesis X-ray, and Genesis Ultrasound, as well as organ-oriented models, including Genesis Brain, Genesis Lung, Genesis Heart, and Genesis Liver.

Cross-domain transfer learning in medical imaging is the Holy Grail. Retrieving a large number of unlabeled images from a PACS system requires an IRB approval, often a long process; the retrieved images must be de-identified; organizing the de-identified images in a way suitable for deep learning is tedious and laborious. Therefore, large quantities of unlabeled datasets may not be readily available to many target domains. Evidenced by our results in Table 3 (BMS), Models Genesis have a great potential for cross-domain transfer learning; particularly, our distortion-based approaches (such as non-linear and local-shuffling) take advantage of relative intensity values (in all modalities) to learn shapes and appearances of various organs. Therefore, as our future work, we will be focusing on methods that generalize well across domains.

5.3. Is any data augmentation suitable as a transformation?

We propose a self-supervised learning framework to learn image representation by discriminating and restoring images undergoing different transformations. One might argue that our image transformations can be interchangeable with existing data augmentation techniques (Gan et al., 2015; Wong et al., 2016; Perez and Wang, 2017; Shorten and Khoshgoftaar, 2019), while we would like to make the distinction between these two concepts clearer. It is critical to assess whether a specific augmentation is practical and feasible for the image restoration task when designing image transformations. Simply introducing data augmentation can make a task ambiguous and lead to degenerate learning. To this end, we choose image transformations based on two principles:

- First, the transformed sub-volume should not be found in the original CT scan. But it is possible to find a transformed sub-volume that has undergone such augmentations as rotation,




Fig. 8. We compare three different losses for the task of image restoration. There is no evidence that the three losses have a decisive impact on the transfer learning results of five target tasks. Note that for this ablation study, all the proxy and target tasks are implemented in PyTorch.

flip, zoom in/out, or translation, as an alternative sub-volume in the original CT scan. In this scenario, without additional spatial information, the model would not be able to “recover” the original sub-volume by seeing the transformed one. As a result, we only elect the augmentations that can be applied to sub-volumes at the pixel level rather than the spatial level.

- Second, a transformation should be applicable for specific image properties. The augmentations that manipulate RGB channels, such as color shift and channel dropping, have little effect on CT/MRI images without the availability of color information. Instead, we promote brightness and contrast into monotonic color curves, resulting in a novel non-linear transformation, explicitly enabling the model to learn intensity distribution from medical images.

After filtering out using the above two principles, the remaining data augmentation techniques are not as many as expected. We have endeavored to produce learning perspective driven transformations rather than inviting any types of data augmentation into our framework. A recent study from Chen et al. (2020) has also discovered a similar phenomenon: carefully designed augmentations are superior to autonomously discovered augmentations. This suggests a criterion of transformations driven by learning perspectives, in capturing a compelling, robust representation for 3D transfer learning in medical imaging.

5.4. Can algorithms autonomously search for transformations?

We follow two principles when designing suitable image transformations for our self-supervised learning framework (see Section 5.3). Potentially, “automated data augmentation” can be considered as an efficient alternative because this line of research seeks to strip researchers from the burden of finding good parameterizations and compositions of transformations manually. Specifically, existing automated augmentation strategies reinforce models to learn an optimal set of augmentation policies by calculating the reward between predictions and image labels. To name a few, Ratner et al. (2017) proposed a method for learning how to parameterize and composite the transformations for automated data augmentation, while preserving class labels or null class for all data points. Dao et al. (2019) introduced a fast kernel alignment metric for augmentation selection. It requires image labels for computing the kernel target alignment (as the reward) between the feature kernel and the label kernel. Cubuk et al. (2019) used reinforcement learning to form an algorithm that autonomously searches for preferred augmentation policies, magnitude, and probability for specific classification tasks, wherein the resultant accuracy of predictions and labels is treated as the reward signal to train the recurrent network controller. Wu et al. (2020) proposed uncertainty-based sampling to select the most effective augmentation, but it is based on the highest loss that is computed between predictions and labels. While the reward is well-defined in the aforementioned works, unfortunately, there is no available metric to

determine the power of image representation directly; hence, no reward is readily established for representation learning. Rather than constrain the representation directly, our paper aims to design an image restoration task to let the model learn generic image representation from 3D medical images. To achieve this, inspired by Vincent et al. (2010), we modify the definition of a good representation into the following: “a good representation is one that can be obtained robustly from a transformed input, and that will be useful for restoring the corresponding original input.” Consequently, mean square error (MSE) between the model’s input and output is defined as the objective function in our framework. However, if we adopt MSE as the reward function, the existing automated augmentation strategies will end up selecting identical-mapping. This is because restoring images without any transformation is expected to give a lower error than restoring those with transformations. Evidenced by Fig. 3, identical-mapping results in a poor image representation. To summarize, the key challenge when employing automated augmentation strategies into our framework is how to define a proper reward for restoring images, and fundamentally, for learning image representation.

5.5. How to assess restoration quality and its relationship to model transferability?

Our transfer learning results in Section 4 suggest that image restoration is a promising task to learn generic 3D image representation. This also means that image restoration quality has an implicit correlation with model transferability to some extent. To assess restoration quality, we compare the Mean Square Error (MSE) loss with other commonly used loss functions for image restoration, such as Mean Absolute Error (MAE) and Structural Similarity Index (SSIM) (Wang et al., 2004). All of them compute the distance between input and output images, while SSIM concentrates more on the restoration quality in terms of structural similarity than MSE and MAE. Since the publicly available 3D SSIM loss was implemented in PyTorch⁷, to make the comparisons fair, we have adapted our five target tasks into PyTorch as well. Fig. 8 shows mixed performances of the five target tasks among the three alternative loss functions. As discussed in Section 5.4, the ideal loss function for representation learning is one that can explicitly determine the power of image representation. However, the three losses explored in this section are implicit, based on the premise that the image restoration quality can indicate a good representation. Further studies with restoration quality assessment and its relationship to model transferability are therefore suggested.

⁷ SSIM loss in 3D: github.com/jinh0park/pytorch-ssim-3D




Fig. 9. [Better viewed on-line and zoomed in for details] Examples of image restoration using Genesis Chest CT. We pass unseen CT images (Column 1) to the pre-trained model, obtaining the restored images (Column 2). The difference between input and output has been shown in Column 3. In most of the normal cases, such as those in Rows 1–2, Genesis Chest CT can perform a fairly reasonable identical-mapping. Meanwhile, for some cases that contain opacity in the lung, as illustrated in Row 3, Genesis Chest CT tends to restore a clearer lung. As a result, the diffuse region is revealed in the difference map automatically. We have zoomed in the region for a better visualization and comparison.

5.6. Could Models Genesis detect infected regions from images autonomously?

As referenced from Section 3.1, Genesis Chest CT has been pre-trained using 623 CT images in the LUNA 2016 dataset. To assess the image restoration quality, we utilize the rest of the 265 CT images from the dataset and present examples in Fig. 9. Specifically, we pass the original CT images to the pre-trained Genesis Chest CT. To visualize the modifications, we have further plotted the difference maps by subtracting the input and output. Since the input images involve no image transformation, most of the restored CT scans (see Rows 1–2) can preserve the texture and structures of the input images, only encountering few changes thanks to the identical-mapping training scheme and the skip connections between encoder and decoder. Nonetheless, we observe some failed cases (see Row 3), especially when the input CT image contains diffuse disease, which appears as an opacity in the lung. Genesis Chest CT happens to “remove” those opaque regions and restore a much clearer lung. This may be due to the fact that the majority of cropped sub-volumes are normal and are being used as ground truth, which empowers the pre-trained model with capabilities of detecting and restoring “novelties” in the CT scans. More specifically, in our work, these novelties include abnormal intensity distribution injected by non-linear transformation, atypical texture and boundary injected by local-shuffling, and discon-

tinuity injected by both inner and outer cutout. Based on the surrounding anatomical structure, the model predicts the opaque area to be air, therefore restoring darker intensity values. This behavior is certainly a “mistake” in terms of image restoration, but it can also be thought of as an attempt to detect diffuse diseases in the lung, which is challenging to annotate due to their unclear boundary. By training an image restoration task, the diseased area will be revealed by simple *subtraction* of the input and output. More importantly, this suggested detection approach requires zero human annotation, neither image-level label nor pixel-level contour, contrasting from the existing weakly supervised disease detection approaches (Zhou et al., 2016; Baumgartner et al., 2018; Cai et al., 2018; Siddiquee et al., 2019).

6. Related work

With the splendid success of deep neural networks, transfer learning (Pan and Yang, 2010; Weiss et al., 2016; Yosinski et al., 2014) has become integral to many applications, especially medical imaging (Greenspan et al., 2016; Litjens et al., 2017; Lu et al., 2017; Shen et al., 2017; Wang et al., 2017a; Zhou et al., 2017; 2019b). This immense popularity of transfer learning is attributed to the learned image representation, which offers convergence speedups and performance gains for most target tasks, in particular, with limited annotated data. In the following sections, we re-

view the works related to supervised and self-supervised representation learning.

6.1. Supervised representation learning

ImageNet contains more than fourteen million images that have been manually annotated to indicate which objects are present in each image; and more than one million of the images have actually been annotated with the bounding boxes of the objects in the image. Pre-training a model on ImageNet and then fine-tuning it on different medical imaging tasks has seen the most practical adoption in medical image analysis (Shin et al., 2016; Tajbakhsh et al., 2016). To classify the common thoracic diseases from ChestX-ray14 dataset, as evidenced in Irvin et al. (2019), nearly all the leading methods (Guan and Huang, 2018; Guendel et al., 2018; Ma et al., 2019; Tang et al., 2018) follow the paradigm of “fine-tuning Models ImageNet” by adopting different architectures, such as ResNet (He et al., 2016) and DenseNet (Huang et al., 2017), along with their pre-trained weights. Other representative medical applications include identifying skin cancer from dermatologist level photographs (Esteva et al., 2017), offering early detection of Alzheimer’s Disease (Ding et al., 2018), and performing effective detection of pulmonary embolism (Tajbakhsh et al., 2019b).

Despite the remarkable transferability of Models ImageNet, pre-trained 2D models offer little benefits towards 3D medical imaging tasks in the most prominent medical modalities (e.g., CT and MRI). To fit this paradigm, 3D imaging tasks have to be reformulated and solved in 2D or 2.5D (Roth et al., 2015; 2014; Tajbakhsh et al., 2015), thus losing rich 3D anatomical information and inevitably compromising the performance. Annotating 3D medical images at the similar scale with ImageNet requires a significant research effort and budget. It is currently not feasible to create annotated datasets comparable to this size for every 3D medical application. Consequently, for lung cancer risk malignancy estimation, Ardila et al. (2019) resorted to incorporate 3D spatial information by using Inflated 3D (I3D) (Carreira and Zisserman, 2017), trained from the Kinetics dataset, as the feature extractor. Evidenced by Table 3, it is not the most favorable choice owing to the large domain gap between the temporal video and medical volume. This limitation has led to the development of model zoo in NiftyNet (Gibson et al., 2018b). However, they were trained with small datasets for specific applications (e.g., brain parcellation and organ segmentation), and were never intended as source models for transfer learning. Our experimental results in Table 3 indicate that NiftyNet models offer limited benefits to the five target medical applications via transfer learning. More recently, Chen et al. (2019b) have pre-trained 3D residual network by jointly segmenting the objects annotated in a collection of eight medical datasets, resulting in MedicalNet for 3D transfer learning. In Table 3, we have examined the pre-trained MedicalNet on five target tasks in comparison with our Models Genesis. As reviewed, each and every aforementioned pre-trained model requires massive, high-quality annotated datasets. However, seldom do we have a perfectly-sized and systematically-labeled dataset to pre-train a deep model in medical imaging, where both data and annotations are expensive to acquire. We overcome the above limitation via self-supervised learning, which allows models to learn image representation from abundant unlabeled medical image data with zero human annotation effort.

6.2. Self-supervised representation learning

Aiming at learning image representation from unlabeled data, self-supervised learning research has recently experienced a surge

in computer vision (Caron et al., 2018; Chen et al., 2019c; Doersch et al., 2015; Goyal et al., 2019; Jing and Tian, 2020; Mahendran et al., 2018; Mundhenk et al., 2018; Noroozi et al., 2018; Noroozi and Favaro, 2016; Pathak et al., 2016; Sayed et al., 2018; Zhang et al., 2016; 2017), but it is a relatively new trend in modern medical imaging. The key challenge for self-supervised learning is identifying a suitable self supervision task, *i.e.*, generating input and output instance pairs from the data. Two of the preliminary studies include predicting the distance and 3D coordinates of two patches randomly sampled from the same brain (Spitzer et al., 2018), identifying whether two scans belong to the same person, and predicting the level of vertebral bodies (Jamaludin et al., 2017). Nevertheless, these two works are incapable of learning representation from “self-supervision” because they demand auxiliary information and specialized data collection such as paired and registered images. By utilizing only the original pixel/voxel information shipped with data, several self-supervised learning schemes have been developed for different medical applications: Ross et al. (2018) adopted colorization as proxy task, wherein color colonoscopy images are converted to gray-scale and then recovered using a conditional Generative Adversarial Network (GAN); Alex et al. (2017) pre-trained a stack of denoising auto-encoders, wherein the self-supervision was created by mapping the patches with the injected noise to the original patches; Chen et al. (2019a) designed image restoration as proxy task, wherein small regions were shuffled within images and then let models learn to restore the original ones; Zhuang et al. (2019) and Zhu et al. (2020) introduced a 3D representation learning proxy task by recovering the rearranged and rotated Rubik’s cube; and finally Tajbakhsh et al. (2019a) individualized self-supervised schemes for a set of target tasks. As seen, the previously discussed self-supervised learning schemes, both in computer vision and medical imaging, are developed individually for specific target tasks, therefore, the generalizability and robustness of the learned image representation have yet to be examined across multiple target tasks. To our knowledge, we are the first to investigate cross-domain self-supervised learning in medical imaging.

6.3. Our previous work

Zhou et al. (2019b) first presented generic autodidactic models for 3D medical imaging, which obtain common image representation that is transferable and generalizable across diseases, organs and modalities, overcoming the scalability issue associated with multiple tasks. This paper extends the preliminary version substantially with the following improvements.

1. We have introduced notations, formulas, and diagrams, as well as detailed methodology descriptions along with their learning objectives, for a succinct framework overview in Section 2.
2. We have extended the brain tumor segmentation experiment using MRI Flair images in Section 3.2, highlighting the transfer learning capabilities of Models Genesis from CT to MRI Flair domains.
3. We have conducted comprehensive ablation studies between the combined scheme and each of the individual learning schemes in Section 4.1, demonstrating that learning from multiple perspectives leads to a more robust target task performance.
4. We have investigated three different random initialization methods for 3D models in Section 4.2, suggesting that initializing with Models Genesis can offer much higher performances and faster convergences.
5. We have examined Models Genesis with the existing pre-trained 3D models on five distinct medical target tasks in Section 4.3, showing that with fewer parameters, Models Gen-

esis surpass all publicly available 3D models in both generalizability and transferability.

6. We have provided experimental results on five target tasks using limited annotated data in [Section 4.4](#), indicating that transfer learning from our Models Genesis can reduce annotation efforts by at least 30%.
7. We have investigated 3D sub-volume based approaches compared with 2D approaches fine-tuning from Models ImageNet using 2D/2.5D representation, underlining the power of pre-trained 3D models in [Section 4.5](#).

7. Conclusion

A key contribution of ours is a collection of *generic source* models, nicknamed Models Genesis, built directly from *unlabeled* 3D imaging data with our novel unified self-supervised method, for generating powerful application-specific *target* models through transfer learning. While the empirical results are strong, surpassing state-of-the-art performances in most of the applications, our goal is to extend our Models Genesis to modality-oriented models, such as Genesis MRI and Genesis Ultrasound, as well as organ-oriented models, such as Genesis Brain and Genesis Heart. We envision that Models Genesis may serve as a primary source of transfer learning for 3D medical imaging applications, in particular, with limited annotated data. To benefit the research community, we make the development of Models Genesis open science, releasing our codes and models to the public. Creating all Models Genesis, an ambitious undertaking, takes a village; therefore, we would like to invite researchers around the world to contribute to this effort, and hope that our collective efforts will lead to the Holy Grail of Models Genesis, all powerful across diseases, organs, and modalities.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRedit authorship contribution statement

Zongwei Zhou: Methodology, Software, Formal analysis, Investigation, Writing - original draft, Writing - review & editing, Visualization. **Vatsal Sodha:** Software, Investigation, Visualization. **Jiaxuan Pang:** Software, Investigation. **Michael B. Gotway:** Resources, Data curation, Funding acquisition. **Jianming Liang:** Conceptualization, Methodology, Formal analysis, Investigation, Resources, Writing - original draft, Writing - review & editing, Supervision, Project administration, Funding acquisition.

Acknowledgments

This research has been supported partially by ASU and Mayo Clinic through a Seed Grant and an Innovation Grant, and partially by the National Institutes of Health (NIH) under Award Number R01HL128785. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH. This work has utilized the GPUs provided partially by the ASU Research Computing and partially by the Extreme Science and Engineering Discovery Environment (XSEDE) funded by the [National Science Foundation](#) (NSF) under grant number [ACI-1548562](#). We thank Z. Guo for implementing Rubik's Cube ([Zhuang et al., 2019](#)) and the 3D version of Jigsaw ([Noroozi and Favaro, 2016](#)) and DeepCluster ([Caron et al., 2018](#)); F. Haghghi and

M. R. Hosseinzadeh Taher for implementing the 3D version of in-painting ([Pathak et al., 2016](#)), patch-shuffling ([Chen et al., 2019a](#)), and working with Z. Guo in evaluating the performance of MedicalNet ([Chen et al., 2019b](#)); M. M. Rahman Siddiquee for examining NiftyNet ([Gibson et al., 2018b](#)) with our Models Genesis; P. Zhang for comparing two additional random initialization methods with our Models Genesis; S. Bajpai for comparing three loss functions of the proxy task; N. Tajbakhsh for revising our conference paper; R. Feng for valuable discussions; and S. Tatapudi for helping improve the writing of this paper. The content of this paper is covered by patents pending.

Appendix A. Implementation details of revised baselines

This work is among the first effort to create a comprehensive benchmark for existing self-supervised learning methods for 3D medical image analysis. We have extended the six most representative self-supervised learning methods into their 3D versions, including De-noising ([Vincent et al., 2010](#)), In-painting ([Pathak et al., 2016](#)), Jigsaw ([Noroozi and Favaro, 2016](#)), and Patch-shuffling ([Chen et al., 2019a](#)). These methods were originally introduced for the purpose of 2D imaging. On the other hand, the most recent 3D self-supervised method ([Zhuang et al., 2019](#)) learns representation by playing a Rubik's cube. We have reimplemented it because their official implementation is not publicly available at the time this paper is written. All of the models are pre-trained using the LUNA 2016 dataset ([Setio et al., 2017](#)) with the same sub-volumes extracted from CT scans as our models (see [Section 3.1](#)). The detailed implementations of the baselines are elaborated in the following sections.

A1. Extended 3D de-noising

In our 3D De-noising, which is inspired by its 2D counterpart ([Vincent et al., 2010](#)), the model is trained to restore the original sub-volume from its transformed one with additive Gaussian noise (randomly sampling $\sigma \in [0, 0.1]$). To correctly restore the original sub-volume, models are required to learn Gabor-like edge detectors when denoising transformed sub-volumes. Following the proposed image restoration training scheme, the auto-encoder network is replaced with a 3D U-Net, wherein the input is a $64 \times 64 \times 32$ sub-volume that has undergone Gaussian noise and the output is the restored sub-volume. The L2 distance between input and output is used as the loss function.

A2. Extended 3D in-painting

In our 3D In-painting, which is inspired by its 2D counterpart ([Pathak et al., 2016](#)), the model is trained to in-paint arbitrary cutout regions based on the rest of the sub-volume. A qualitative illustration of the image in-painting task is shown in the right panel of [Fig. A.11\(a\)](#). To correctly predict missing regions, networks are required to learn local continuities of organs in medical images via interpolation. Unlike the original in-painting, the adversarial loss and discriminator are excluded from our implementation of the 3D version because our primary goal is to empower models with generic representation, rather than generating sharper and realistic sub-volumes. The generator is a 3D U-Net, consisting of an encoder and a decoder. The input of the encoder is a $64 \times 64 \times 32$ sub-volume that needs to be in-painted. Their decoder works differently than our inner-cutout because it predicts the missing region only, and therefore, the loss is just computed on the cutout region—an ablation study on the loss has been further presented in [Appendix C.2](#).

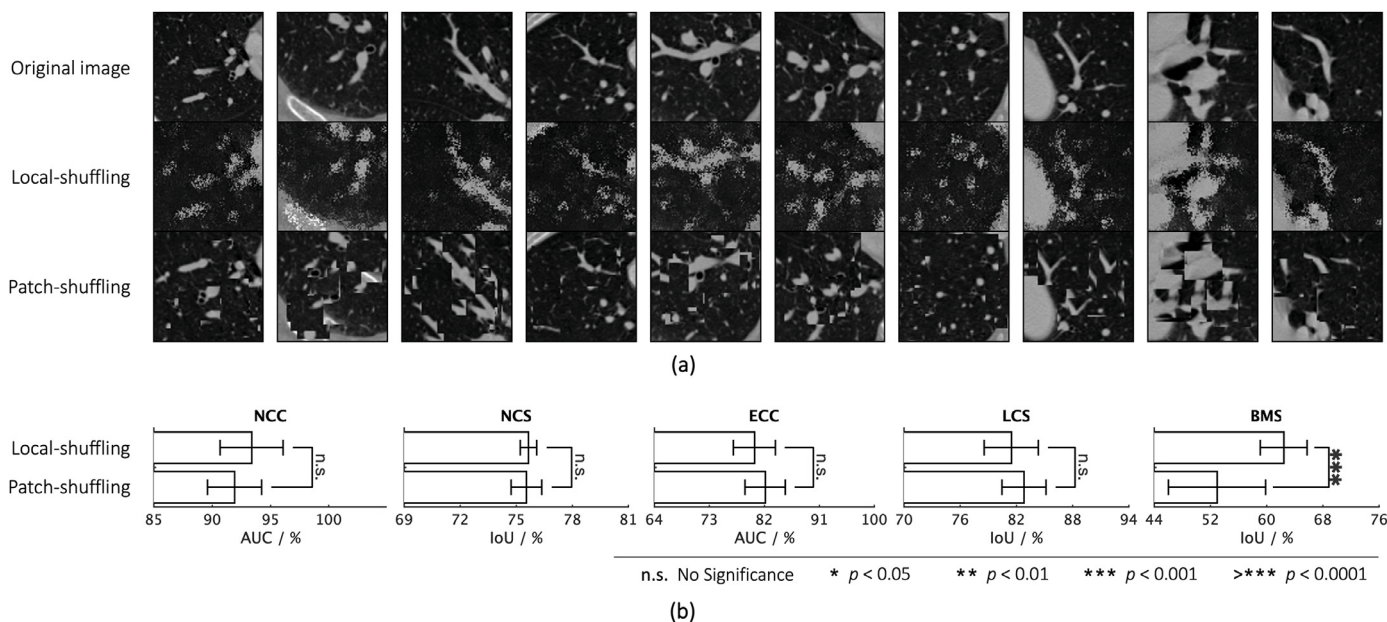


Fig. A.10. A direct comparison between global patch shuffling (Chen et al., 2019a) and our local pixel shuffling. (a) illustrates ten example images undergone local-shuffling and patch-shuffling independently. As seen, the overall anatomical structure such as individual organs, blood vessels, lymph nodes, and other soft tissue structures are preserved in the transformed image through local-shuffling. (b) presents the performance on five target tasks, showing that models pre-trained by our local-shuffling noticeably outperform those pre-trained by patch-shuffling for cross-domain transfer learning (BMS).

A3. Extended 3D jigsaw

In our 3D Jigsaw, which is inspired by its 2D counterpart (Noroozi and Favaro, 2016), we utilize the implementation by Taleb et al. (2020)⁸, wherein the puzzles are created by sampling a $3 \times 3 \times 3$ grid of 3D patches. Then, these patches are shuffled according to an arbitrary permutation, selected from a set of predefined permutations. This set with size $P = 100$ is chosen out of the $(3 \times 3 \times 3)!$ possible permutations, by following the Hamming distance based algorithm, and each permutation is assigned an index. As a result, the problem is cast as a P -way classification task, i.e., the model is trained to recognize the applied permutation index, allowing us to solve the 3D puzzles efficiently. We build the classification model by taking the encoder of 3D U-Net and appending a sequence of fc layers. In the implementation, we minimize the cross-entropy loss of the list of extracted puzzles.

A4. Extended 3D patch-shuffling

In our 3D Patch-shuffling, which is inspired by its 2D counterpart (Chen et al., 2019a), the model learns image representation by restoring the image context. Given a sub-volume, we randomly select two isolated small 3D patches and swap their context. We set the length, width, and height of the 3D patch to be proportional to those in the entire sub-volume by 25% to 50%. Repeating this process for $T = 10$ times can generate the transformed sub-volume (see examples in Fig. A.10(a)). The model is trained to restore the original sub-volume, where L2 distance between input and output is used as the loss function. To process volumetric input and ensure a fair comparison with other baselines, we replace their U-Net with 3D U-Net architecture, where the encoder and decoder serve as analysis and restoration parts, respectively.

⁸ Self-Supervised 3D Tasks: github.com/HealthML/self-supervised-3d-tasks

A5. Extended 3D deepcluster

In our 3D DeepCluster, which is inspired by its 2D counterpart (Caron et al., 2018), we iteratively cluster deep features extracted from sub-volumes by k -means and use the subsequent assignments as supervision to update the weights of the model. Through clustering, the model can obtain useful general-purpose visual features, requiring little domain knowledge and no specific signal from the inputs. We replaced original AlexNet/VGG architecture with the encoder of 3D U-Net to process 3D input sub-volumes. The number of clusters that works best for 2D tasks may not be a good choice for 3D tasks. To ensure a fair comparison, we extensively tune this hyper-parameter in $\{10, 20, 40, 80, 160, 320\}$ and finally set to 260 from the narrowed down search space of $\{240, 260, 280\}$. Unlike ImageNet models for 2D imaging tasks, there is no available pre-trained 3D feature extractor for medical imaging tasks; therefore, we randomly initialize the model weights at the beginning. Our Models Genesis, the first generic 3D pre-trained models, could potentially be used as the 3D feature extractor and co-trained with 3D DeepCluster.

A6. Rubik's cube

We implement Rubik's Cube with respect to Zhuang et al. (2019), which consists of cube rearrangement and cube rotation. Like playing a Rubik's cube, this proxy task enforces models to learn translational and rotational invariant features from raw 3D data. Given a sub-volume, we partition it into a $2 \times 2 \times 2$ grid of cubes. In addition to predicting orders (3D Jigsaw), this proxy task permutes the cubes with random rotations, forcing models to predict the orientation. Following the original paper, we limit the directions for cube rotation, i.e., only allowing 180° horizontal and vertical rotations, to reduce the complexity of the task. The eight cubes are then fed into a Siamese network with eight branches sharing the

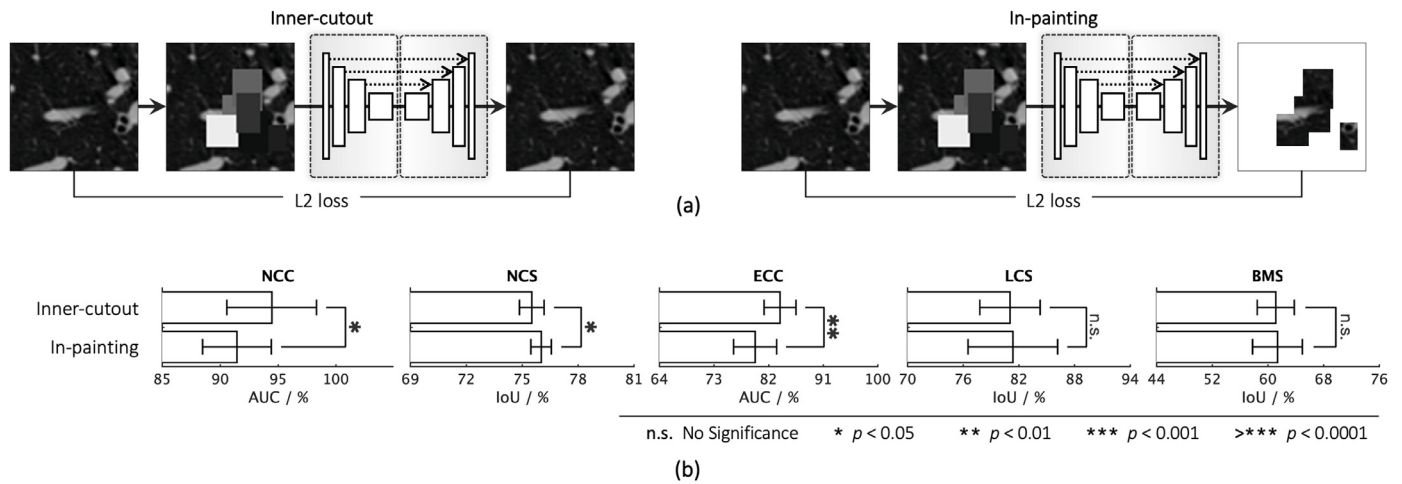


Fig. A.11. A direct comparison between image in-painting (Pathak et al., 2016) and our inner-cutout. (a) contrasts our inner-cutout with in-painting, wherein the model in the former scheme computes loss on the entire image and the model in the latter scheme computes loss only for the cutout area. (b) presents the performance on five target tasks, showing that inner-cutout is better suited for target classification tasks (e.g., NCC and ECC), while in-painting is more helpful for target segmentation tasks (e.g., NCS, LCS, and BMS).

same weight to extract features. The feature maps from the last fully-connected or convolution layer of all branches are concatenated and given as input to the fully-connected layer of separate tasks, *i.e.*, cube ordering and orienting, which are supervised by permutation loss and rotation loss, respectively, with equal weights.

Appendix B. Configurations of publicly available models

For publicly available models, we do not re-train their proxy tasks and instead simply endeavor to find the best hyper-parameters for each of them in target tasks. We compare them with our Models Genesis in a user perspective, which might seem to be unfair in a research perspective because many variables are asymmetric among the competitors, such as programming platform, model architecture, number of parameters, etc. However, the goal of this section is to experiment with existing ready-to-use pre-trained models under different medical tasks; therefore, we presume that all of the publicly available models and their configurations have been carefully composed to the optimal setting.

B1. NiftyNet

We examine the effectiveness of fine-tuning from NiftyNet in five target tasks. We should note that NiftyNet is not initially designed for transfer learning but is one of the few publicly available supervised pre-trained 3D models. The model from Gibson et al. (2018a) has been considered as the baseline in our experiments because it has also been pre-trained on the chest region in CT modality and applied an encoder-decoder architecture that is similar to our work. We directly adopt the pre-trained weights of the dense V-Net architecture provided by NiftyNet, so it carries a smaller number of parameters than our 3D U-Net (2.60M vs. 16.32M). For target classification tasks, we use the dense V-Net encoder by appending a sequence of *fc* layers; for target segmentation tasks, we use the entire dense V-Net. Since NiftyNet is developed in Tensorflow, all five target tasks are re-implemented using their build-in configuration. For each target task, we have tuned hyper-parameters (e.g., learning rate and opti-

mizer) and applied extensive data augmentations (e.g., rotation and scaling).

B2. Inflated 3D

We download the Inflated 3D (I3D) model pre-trained from Flow streams in the Kinetics dataset (Hara et al., 2018) and fine-tune it on our five target tasks. The input sub-volume is copied into two channels to align with the required input shape. For target classification tasks, we take the pre-trained I3D and append a sequence of randomly initialized fully-connected layers. For target segmentation tasks, we take the pre-trained I3D as the encoder and expand a decoder to predict the segmentation map, resulting in a U-Net like architecture. The decoder is the same as that implemented in our 3D U-Net, consisting of up-sampling layers followed by a sequence of convolutional layers, batch normalization, and ReLU activation. Besides, four skip connections are built between the encoder and decoder, wherein feature maps before each pooling layer in the encoder are concatenated with same-scale feature maps in the decoder. All of the layers in the model are trainable during transfer learning. Adam method (Kinga and Adam, 2015) with a learning rate of $1e - 4$ is used for optimization.

B3. MedicalNet

We download MedicalNet models (Chen et al., 2019b) that have been pre-trained on eight publicly available 3D segmentation datasets. ResNet-50 and ResNet-101 backbones are chosen because they are reported by Chen et al. (2019b) as the most compelling backbones for target segmentation and classification tasks, respectively. Like I3D, we append a decoder at the end of the pre-trained encoder, randomly initialize its weights, and link the encoder with the decoder using skip connections. Owing to the 3D ResNet backbones, the resultant segmentation network for MedicalNet is much heavier than our 3D U-Net. To be consistent with the original programming platform of MedicalNet, we re-implement all five target tasks in PyTorch, using the same data separation and augmentation. We report the highest results achieved by any of the two backbones in Table 3.

Appendix C. Ablation experiments

C1. Local pixel shuffling vs. global patch shuffling

In the main paper, we have reported results of patch-shuffling (Chen et al., 2019a) as a baseline in Table 3 and our local-shuffling in Fig. 3. To underline the value of preserving local and global structural consistency in the proxy task, we provide an explicit comparison between the two counterparts in Fig. A.10, arriving at three findings:

- Global patch shuffling preserves local information while distorting global structure; local pixel shuffling maintains global structure but loses local details.
- For same-domain transfer learning (e.g., pre-training and fine-tuning in CT images), global-shuffling and local-shuffling reveal no significant difference in terms of target task performance. Note that local-shuffling is preferable when recognizing small objects in target tasks (e.g., pulmonary nodule and embolism), whereas patch-shuffling is beneficial for large objects (e.g., brain tumor and liver).
- For cross-domain transfer learning (e.g., pre-training in CT and fine-tuning in MRI images), models pre-trained by our local-shuffling noticeably outperform those pre-trained by patch-shuffling.




Fig. C.12. We extensively search for the optimal size of cutout regions spanning from 0% to 90%, incremented by 10%. The points plotted within the red shade denote no significant difference ($p > .05$) from the pinnacle from the curve. The horizontal red and gray lines refer to the performances achieved by Models Genesis and learning from scratch, respectively. This ablation study reveals that cutting 20%–40% of regions out could produce the most robust performance of target tasks. As a result, in our implementation, we cutout around 25% regions from each sub-volume. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

C2. Compute loss on cutouts vs. entire images

The results of our ablation study for in-painting and inner-cutout on five target tasks are presented in Fig. A.11. We set all the hyper-parameters the same except for one factor: where to compute MSE loss, only cutout areas or the entire image. In general, there is a marginal difference in target segmentation tasks, but inner-cutout is superior to in-painting in target classification tasks. These results are in line with our hypothesis in Section 3.1: the model must distinguish original versus transformed parts within the image, preserving the context if it is original and, otherwise, in-painting the context. Seemingly, in-painting that only computes loss on cutouts can fail to learn comprehensive representation as it is unable to leverage advancements from both ends.

C3. Masked area size in outer-cutout

When applying cutout transformations to our self-supervised learning framework, we have one hyper-parameter to evaluate, i.e., the size of cutout regions. Intuitively, it can influence the difficulty of the image restoration task. To explore the impact of this parameter on the performance of target tasks, we have conducted an ablation study to extensively search for the optimal value, spanning from 0% to 90%, incremented by intervals of 10%. Fig. C.12 shows the performance of all five target tasks under different settings, suggesting that outer-cutout is robust to hyper-parameter changes to some extent. This finding is also consistent with that recommended in the original in-painting paper, where Pathak et al. (2016) removed a number of smaller possibly overlapping masks, covering up to 1/4 of the image. Altogether, we finally cutout less than 1/4 of the entire sub-volume in both outer and inner cutout implementations.

Appendix D. Qualitative assessment of image restoration

Since there is no such metric to directly determine the power of image representation, rather than constrain the representation, our paper aims to design an image restoration task to let the model learn generic image representation from 3D medical images. In doing so, as seen in Section 5.4, we have modified the definition of a good representation. As presented in Section 3.1, Genesis CT and Genesis X-ray are pre-trained on LUNA 2016 (Setio et al., 2017) and ChestX-ray8 (Wang et al., 2017b), respectively, using a series of self-supervised learning schemes with different image transformations. In this section, we have (1) illustrated more examples of our four individual transformations (i.e., non-linear, local-shuffling, outer-cutout, and inner-cutout) in Fig. D.13; (2) evaluated the power of the pre-trained model by assessing restoration quality on previously *unseen* patients' images not only from the LUNA 2016 dataset (see Fig. D.14), but also from different modalities, covering CT, X-ray, and MRI (see Fig. D.15). The qualitative assessment shows that our pre-trained model is not merely overfitting on anatomical patterns in specific patients, but indeed can be robustly used for restoring images, thus can be generalized to many target tasks.

To assess the image restoration quality at the time of inference, we pass the transformed images to the models that have been trained with different self-supervised learning schemes, including four individual and one combined schemes. In our visualization, the input images have undergone four individual transformations as well as eight different combined transformations, including the identity mapping (i.e., no transformation). As shown in Fig. D.14, the combined scheme can restore the unseen image by handling a variety of transformations (framed in red), whereas the models

trained with the individual scheme can only restore unseen images that have undergone the same transformation that they were trained on (framed in green). This qualitative observation is consistent with our experimental finding in Section 4.1: the combined learning scheme achieves superior and more robust results over the individual scheme in transfer learning.

In Fig. D.15, we have further provided a qualitative assessment of image restoration quality by Genesis CT and Genesis X-ray, across medical imaging modalities. In our visualization, the input

images are selected from four different medical modalities, covering X-ray, CT, Ultrasound, and MRI. It is clear from the figure that even though the models are only trained on single image modality, they can largely maintain the texture and structures during restoration, not only within the same modality but also across different ones. These observations suggest that Models Genesis are of great potential in transferring learned image representation across diseases, organs, datasets, and modalities.

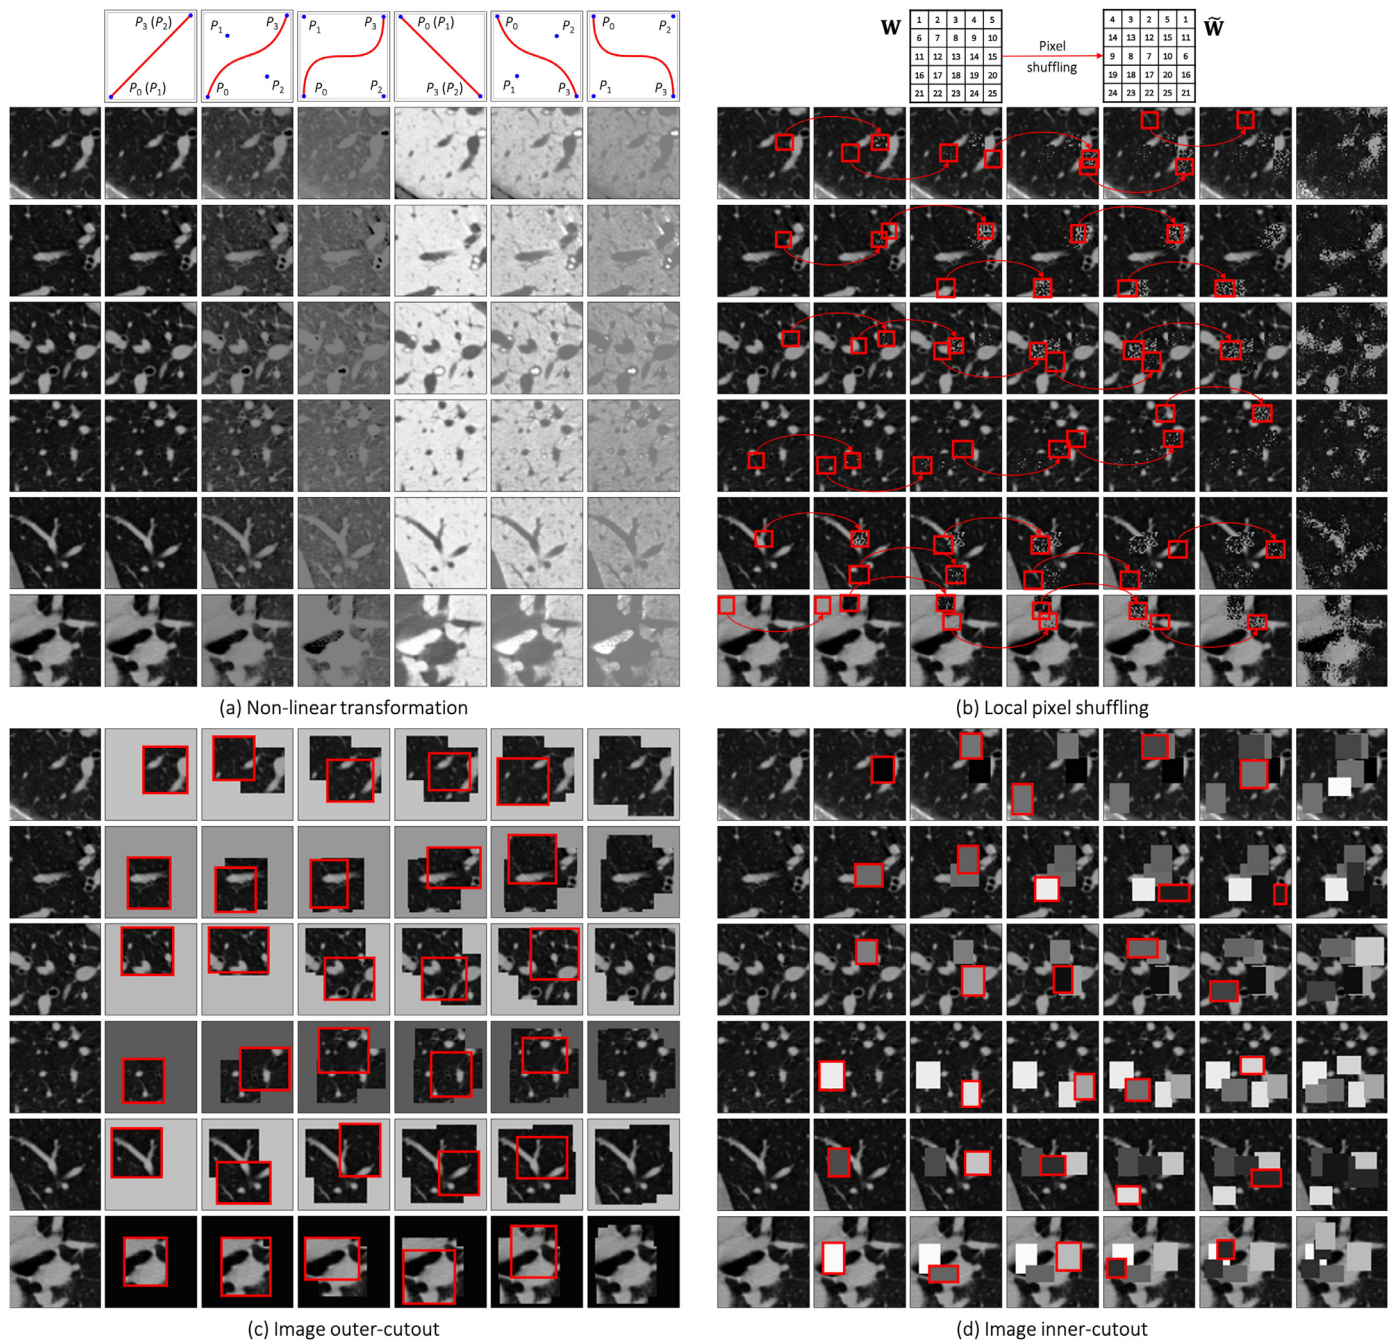


Fig. D.13. Illustration of the proposed four image transformations. For simplicity and clarity, we illustrate the transformation on a 2D CT slice, but our Genesis Chest CT is trained using 3D sub-volumes directly, transformed in a 3D manner; our 3D image transformations, with an exception of non-linear transformation, cannot be approximated in 2D. For ease of understanding, in (a) non-linear transformation, we have displayed an image undergoing different translating functions in Columns 2–7; in (b) local-shuffling, (c) outer-cutout, and (d) inner-cutout transformation, we have illustrated each of the processes step by step in Columns 2–6, where the first and last columns denote the original images and the final transformed images, respectively. In local-shuffling, a different window (b) is automatically generated and used in each step.

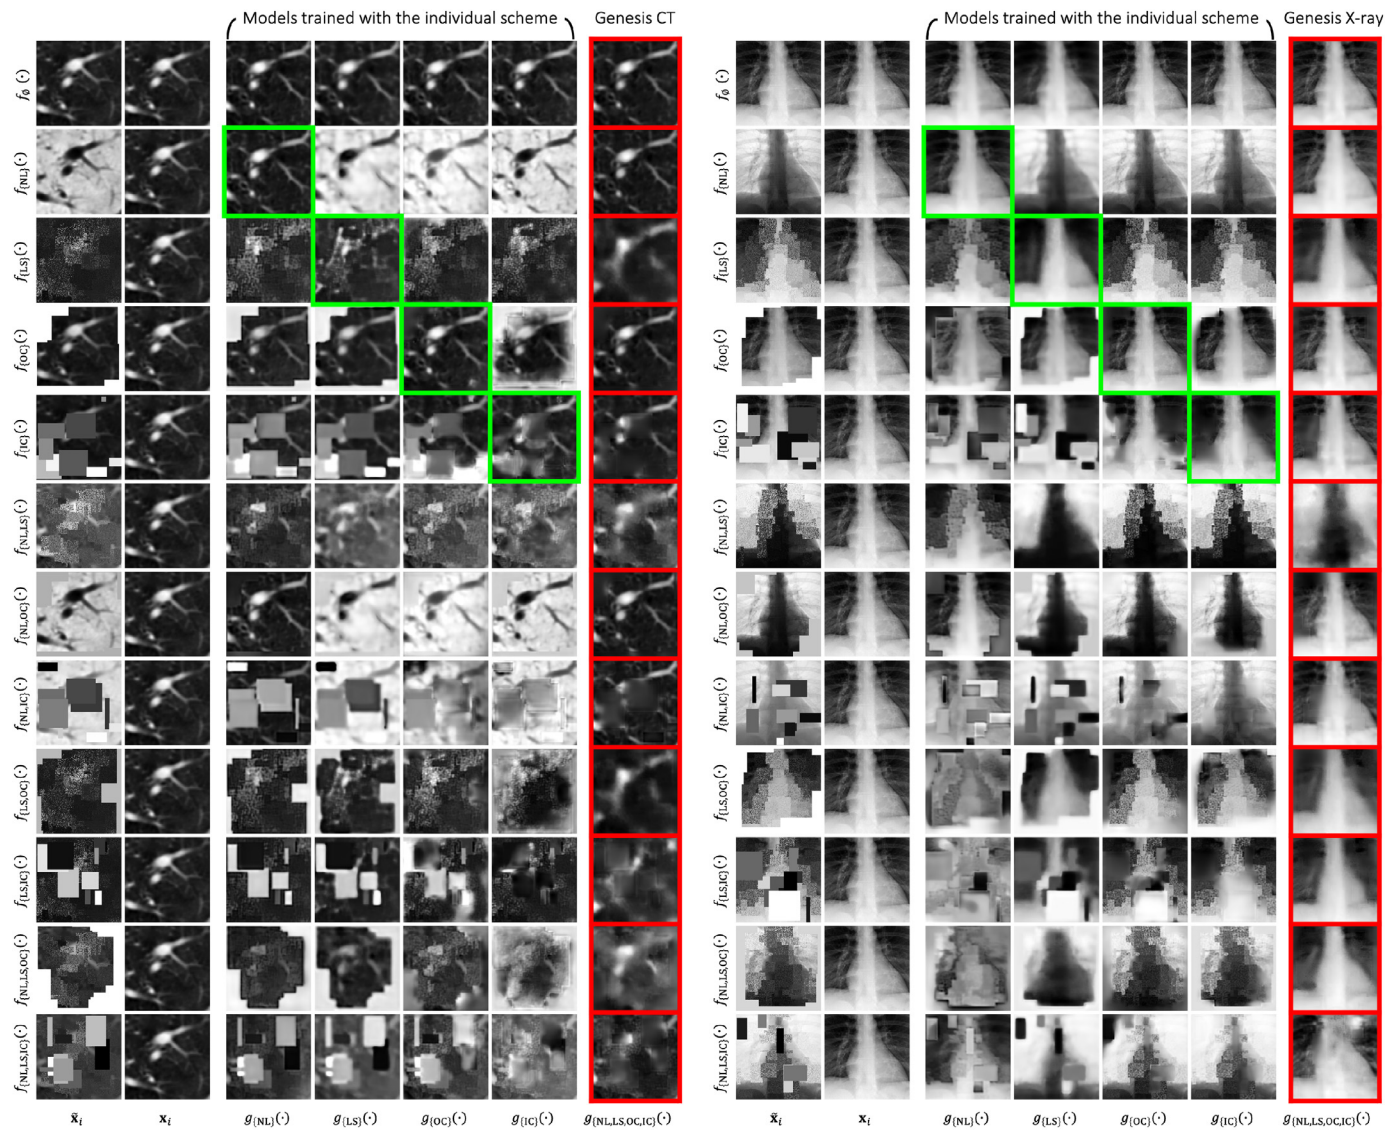


Fig. D.14. The left and right panels show the qualitative assessment of image restoration quality using Genesis CT and Genesis X-ray, respectively. These models are trained with different training schemes, including four individual schemes (Columns 3–6) and a combined scheme (Column 7). As discussed in Fig. 1, each original image x_i can possibly undergo twelve different transformations. We test the models with all these possible twelve transformed images \tilde{x}_i . We specify types of the image transformation $f(\cdot)$ for each row and the training scheme $g(\cdot)$ for each column. First of all, it can be seen that the models trained with individual schemes can restore previously unseen images that have undergone the same transformation very well (framed in green), but fail to handle other transformations. Taking non-linear transformation $f_{NL}(\cdot)$ as an example, any individual training scheme besides non-linear transformation itself cannot invert the pixel intensity from transformed whitish to the original blackish. As expected, the model trained with the combined scheme successfully restores original images from various transformations (framed in red). Second, the model trained with the combined scheme shows it is superior to other models even if they are trained with and tested on the same transformation. For example, in the local-shuffling case $f_{LS}(\cdot)$, the image recovered from the local-shuffling pre-trained model $g_{LS}(\cdot)$ is noisy and lacks texture. However, the model trained with the combined scheme $g_{NLLSOC}(\cdot)$ generates an image with more underlying structures, which demonstrates that learning with augmented tasks can even improve the performance on each of the individual tasks. Third, the model trained with the combined scheme significantly outperforms models trained with individual training schemes when restoring images that have undergone seven different combined transformations (Rows 6–12). For example, the model trained with non-linear transformation $g_{NL}(\cdot)$ can only recover the intensity distribution in the transformed image undergone $f_{NLIC}(\cdot)$ but leaves the inner cutouts unchanged. These observations suggest that the model trained with the proposed unified self-supervised learning framework can successfully learn general anatomical structures and yield promising transferability on different target tasks. The quality assessment of image restoration further confirms our experimental observation, provided in Section 4.1, that the combined learning scheme exceeds each individual in transfer learning. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)




Fig. D.15. The left and right panels visualize the qualitative assessment of image restoration quality by Genesis CT and Genesis X-ray, respectively, across medical imaging modalities. For testing, we use the pre-trained model to directly restore images from LUNA 2016 (CT) (Setio et al., 2017), ChestX-ray8 (X-ray) (Wang et al., 2017b), CIMT (Ultrasound) (Hurst et al., 2010; Zhou et al., 2019a), and BraTS (MRI) (Menze et al., 2015). Though the models are only trained on single image modality, they can largely maintain the texture and structures during restoration not only within the same modality (framed in red), but also across different modalities. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

References

- Alex, V., Vaidhya, K., Thirunavukkarasu, S., Kesavadas, C., Krishnamurthi, G., 2017. Semisupervised learning using denoising autoencoders for brain lesion detection and segmentation. *J. Med. Imaging* 4 (4), 041311.
- Ardila, D., Kiraly, A.P., Bharadwaj, S., Choi, B., Reicher, J.J., Peng, L., Tse, D., Etemadi, M., Ye, W., Corrado, G., et al., 2019. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. *Nat. Med.* 25 (6), 954–961.
- Aresta, G., Jacobs, C., Araújo, T., Cunha, A., Ramos, I., van Ginneken, B., Campilho, A., 2019. Iw-net: an automatic and minimalistic interactive lung nodule segmentation deep network. *Sci. Rep.* 9 (1), 1–9.
- Armato III, S.G., McLennan, G., Bidaut, L., McNitt-Gray, M.F., Meyer, C.R., Reeves, A.P., Zhao, B., Aberle, D.R., Henschke, C.I., Hoffman, E.A., et al., 2011. The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Med. Phys.* 38 (2), 915–931.
- Bakas, S., Reyes, M., Jakab, A., Bauer, S., Rempfler, M., Crimi, A., Shinohara, R.T., Berger, C., Ha, S.M., Rozycki, M., et al., 2018. Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. *arXiv preprint arXiv:1811.02629*.
- Baumgartner, C.F., Koch, L.M., Can Tezcan, K., Xi Ang, J., Konukoglu, E., 2018. Visual feature attribution using wasserstein gans. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8309–8319.
- Ben-Cohen, A., Diamant, I., Klang, E., Amitai, M., Greenspan, H., 2016. Fully Convolutional Network for Liver Segmentation and Lesions Detection. In: *Deep learning and data labeling for medical applications*. Springer, pp. 77–85.
- Bilic, P., Christ, P.F., Vorontsov, E., Chlebus, G., Chen, H., Dou, Q., Fu, C.-W., Han, X., Heng, P.-A., Hesser, J., et al., 2019. The liver tumor segmentation benchmark (lits). *arXiv preprint arXiv:1901.04056*.
- Buzug, T.M., 2011. *Computed Tomography*. In: *Springer Handbook of Medical Technology*. Springer, pp. 311–342.
- Cai, J., Lu, L., Harrison, A.P., Shi, X., Chen, P., Yang, L., 2018. Iterative attention mining for weakly supervised thoracic disease pattern localization in chest x-rays. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 589–598.
- Caron, M., Bojanowski, P., Joulin, A., Douze, M., 2018. Deep clustering for unsupervised learning of visual features. In: *Proceedings of the European Conference on Computer Vision*, pp. 132–149.
- Caron, M., Bojanowski, P., Mairal, J., Joulin, A., 2019. Unsupervised pre-training of image features on non-curved data. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2959–2968.
- Carreira, J., Zisserman, A., 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6299–6308.
- Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., Rueckert, D., 2019. Self-supervised learning for medical image analysis using image context restoration. *Med. Image Anal.* 58, 101539.
- Chen, S., Ma, K., Zheng, Y., 2019. Med3d: transfer learning for 3d medical image analysis. *arXiv preprint arXiv:1904.00625*.
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G., 2020. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*.
- Chen, T., Zhai, X., Ritter, M., Lucic, M., Houlsby, N., 2019. Self-supervised gans via auxiliary rotation loss. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 12154–12163.
- Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V., 2019. Autoaugment: Learning augmentation strategies from data. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1113–1123.
- Dao, T., Gu, A., Ratner, A.J., Smith, V., De Sa, C., Ré, C., 2019. A kernel theory of modern data augmentation. *Proceedings of machine learning research* 97, 1528.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, pp. 248–255.
- Ding, Y., Sohn, J.H., Kawczynski, M.G., Trivedi, H., Harnish, R., Jenkins, N.W., Lituev, D., Copeland, T.P., Aboian, M.S., Mari Aparici, C., et al., 2018. A deep learning model to predict a diagnosis of alzheimer disease by using 18f-fdg pet of the brain. *Radiology* 290 (2), 456–464.
- Doersch, C., Gupta, A., Efros, A.A., 2015. Unsupervised visual representation learning by context prediction. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1422–1430.
- Doersch, C., Zisserman, A., 2017. Multi-task self-supervised visual learning. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2051–2060.
- Dosovitskiy, A., Fischer, P., Springenberg, J.T., Riedmiller, M., Brox, T., 2015. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 38 (9), 1734–1747.
- Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S., 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542 (7639), 115.
- Forbes, G.B., 2012. *Human body composition: Growth, aging, nutrition, and activity*. Springer Science & Business Media.
- Gan, Z., Henao, R., Carlson, D., Carin, L., 2015. Learning deep sigmoid belief networks with data augmentation. In: *Artificial Intelligence and Statistics*, pp. 268–276.
- Gibson, E., Giganti, F., Hu, Y., Bonmati, E., Bandula, S., Gurusamy, K., Davidson, B., Pereira, S.P., Clarkson, M.J., Barratt, D.C., 2018. Automatic multi-organ segmentation on abdominal ct with dense v-networks. *IEEE Trans. Med. Imaging* 37 (8), 1822–1834.
- Gibson, E., Li, W., Sudre, C., Fidon, L., Shaker, D.I., Wang, G., Eaton-Rosen, Z., Gray, R., Doel, T., Hu, Y., et al., 2018. Niftnet: a deep-learning platform for medical imaging. *Comput. Methods Programs Biomed.* 158, 113–122.
- Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks. In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pp. 249–256.
- Goyal, P., Mahajan, D., Gupta, A., Misra, I., 2019. Scaling and benchmarking self-supervised visual representation learning. *arXiv preprint arXiv:1905.01235*.
- Graham, B., 2014. Fractional max-pooling. *arXiv preprint arXiv:1412.6071*.
- Greenspan, H., van Ginneken, B., Summers, R.M., 2016. Guest editorial deep learning in medical imaging: overview and future promise of an exciting new technique. *IEEE Trans Med Imaging* 35 (5), 1153–1159.
- Guan, Q., Huang, Y., 2018. Multi-label chest x-ray image classification via category-wise residual attention learning. *Pattern Recognit. Lett.*
- Guendel, S., Grbic, S., Georgescu, B., Liu, S., Maier, A., Comaniciu, D., 2018. Learning to recognize abnormalities in chest x-rays with location-aware dense networks. In: *Iberoamerican Congress on Pattern Recognition*. Springer, pp. 757–765.
- Hara, K., Kataoka, H., Satoh, Y., 2018. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6546–6555.
- He, K., Zhang, X., Ren, S., Sun, J., 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1026–1034.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778.
- Hendrycks, D., Mazeika, M., Kadavath, S., Song, D., 2019. Using self-supervised learning can improve model robustness and uncertainty. In: *Advances in Neural Information Processing Systems*, pp. 15637–15648.
- Huang, G., Liu, Z., Weinberger, K.Q., van der Maaten, L., 2017. Densely connected convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1, p. 3.
- Hurst, R.T., Burke, R.F., Wissner, E., Roberts, A., Kendall, C.B., Lester, S.J., Somers, V., Goldman, M.E., Wu, Q., Khandheria, B., 2010. Incidence of subclinical atherosclerosis as a marker of cardiovascular risk in retired professional football players. *Am. J. Cardiol.* 105 (8), 1107–1111.
- Iizuka, S., Simo-Serra, E., Ishikawa, H., 2017. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)* 36 (4), 107.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In: *Bach, F., Blei, D. (Eds.), Proceedings of the 32nd International Conference on Machine Learning*. PMLR, Lille, France, pp. 448–456.
- Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoob, B., Ball, R., Shpanskaya, K., et al., 2019. Chexpert: a large chest radiograph dataset with uncertainty labels and expert comparison. *arXiv preprint arXiv:1901.07031*.
- Jamaludin, A., Kadir, T., Zisserman, A., 2017. Self-supervised Learning for Spinal Mris. In: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, pp. 294–302.
- Jing, L., Tian, Y., 2020. Self-supervised visual feature learning with deep neural networks: a survey. *IEEE Trans. Pattern Anal. Mach. Intell.*
- Kang, G., Dong, X., Zheng, L., Yang, Y., 2017. Patchshuffle regularization. *arXiv preprint arXiv:1707.07103*.
- Kinga, D., Adam, J.B., 2015. Adam: A method for stochastic optimization. In: *International Conference on Learning Representations (ICLR)*, 5.
- Kolesnikov, A., Zhai, X., Beyer, L., 2019. Revisiting self-supervised visual representation learning. In: *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 1920–1929.
- Liang, J., Bi, J., 2007. Computer aided detection of pulmonary embolism with toboggaming and multiple instance classification in ct pulmonary angiography. In: *Biennial International Conference on Information Processing in Medical Imaging*. Springer, pp. 630–641.
- Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., Van Der Laak, J.A., Van Ginneken, B., Sánchez, C.I., 2017. A survey on deep learning in medical image analysis. *Med. Image Anal.* 42, 60–88.
- LiTS, 2017. *Results of all submissions for liver segmentation*.
- Lu, L., Zheng, Y., Carneiro, G., Yang, L., 2017. *Deep learning and convolutional neural networks for medical image computing: Precision medicine, high performance and large-scale datasets*. Springer.
- LUNA, 2016. *Results of all submissions for nodule false positive reduction*.
- Ma, Y., Zhou, Q., Chen, X., Lu, H., Zhao, Y., 2019. Multi-attention network for thoracic disease classification and localization. In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, pp. 1378–1382.
- Mahendran, A., Thewlis, J., Vedaldi, A., 2018. Cross pixel optical-flow similarity for self-supervised learning. In: *Asian Conference on Computer Vision*. Springer, pp. 99–116.
- Menze, B.H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., et al., 2015. The multimodal brain tumor image segmentation benchmark (brats). *IEEE Trans. Med. Imaging* 34 (10), 1993.
- Mortenson, M.E., 1999. *Mathematics for computer graphics applications*. Industrial Press Inc..

- Mundhenk, T.N., Ho, D., Chen, B.Y., 2018. Improvements to context based self-supervised learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 9339–9348.
- NLST, 2011. Reduced lung-cancer mortality with low-dose computed tomographic screening. *N top N. Engl. J. Med.* 365 (5), 395–409.
- Noroozi, M., Favaro, P., 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In: European Conference on Computer Vision. Springer, pp. 69–84.
- Noroozi, M., Vinjimoor, A., Favaro, P., Pirsiavash, H., 2018. Boosting self-supervised learning via knowledge transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 9359–9367.
- Pan, S.J., Yang, Q., 2010. A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* 22 (10), 1345–1359.
- Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A., 2016. Context encoders: Feature learning by inpainting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2536–2544.
- Perez, L., Wang, J., 2017. The effectiveness of data augmentation in image classification using deep learning. arXiv preprint arXiv:1712.04621.
- Prasoon, A., Petersen, K., Igel, C., Lauze, F., Dam, E., Nielsen, M., 2013. Deep feature learning for knee cartilage segmentation using a triplanar convolutional neural network. In: International conference on medical image computing and computer-assisted intervention. Springer, pp. 246–253.
- Ratner, A.J., Ehrenberg, H., Hussain, Z., Dunmon, J., Ré, C., 2017. Learning to compose domain-specific transformations for data augmentation. In: Advances in neural information processing systems, pp. 3236–3246.
- Rawat, W., Wang, Z., 2017. Deep convolutional neural networks for image classification: a comprehensive review. *Neural. Comput.* 29 (9), 2352–2449.
- Ross, T., Zimmerer, D., Vemuri, A., Isensee, F., Wiesenfarth, M., Bodenstedt, S., Both, F., Kessler, P., Wagner, M., Müller, B., et al., 2018. Exploiting the potential of unlabeled endoscopic video data with self-supervised learning. *Int. J. Comput. Assist. Radiol. Surg.* 13 (6), 925–933.
- Roth, H.R., Lu, L., Liu, J., Yao, J., Seff, A., Cherry, K., Kim, L., Summers, R.M., 2015. Improving computer-aided detection using convolutional neural networks and random view aggregation. *IEEE Trans. Med. Imaging* 35 (5), 1170–1181.
- Roth, H.R., Lu, L., Seff, A., Cherry, K.M., Hoffman, J., Wang, S., Liu, J., Turkbey, E., Summers, R.M., 2014. A new 2.5 d representation for lymph node detection using random sets of deep convolutional neural network observations. In: International conference on medical image computing and computer-assisted intervention. Springer, pp. 520–527.
- Sayed, N., Brattoli, B., Ommer, B., 2018. Cross and learn: Cross-modal self-supervision. In: German Conference on Pattern Recognition. Springer, pp. 228–243.
- Setio, A.A.A., Ciompi, F., Litjens, G., Gerke, P., Jacobs, C., Van Riel, S.J., Wille, M.M.W., Naqibullah, M., Sánchez, C.I., van Ginneken, B., 2016. Pulmonary nodule detection in ct images: false positive reduction using multi-view convolutional networks. *IEEE Trans. Med. Imaging* 35 (5), 1160–1169.
- Setio, A.A.A., Traverso, A., De Bel, T., Berens, M.S., van den Bogaard, C., Cerello, P., Chen, H., Dou, Q., Fantacci, M.E., Geurts, B., et al., 2017. Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the luna16 challenge. *Med. Image Anal.* 42, 1–13.
- Shen, D., Wu, G., Suk, H.-I., 2017. Deep learning in medical image analysis. *Annu Rev Biomed Eng* 19, 221–248.
- Shin, H.-C., Roth, H.R., Gao, M., Lu, L., Xu, Z., Nogues, I., Yao, J., Mollura, D., Summers, R.M., 2016. Deep convolutional neural networks for computer-aided detection: CNN architectures, dataset characteristics and transfer learning. *IEEE Trans. Med. Imaging* 35 (5), 1285–1298.
- Shorten, C., Khoshgoftaar, T.M., 2019. A survey on image data augmentation for deep learning. *J. Big Data* 6 (1), 60.
- Siddiquee, M.M.R., Zhou, Z., Tajbakhsh, N., Feng, R., Gotway, M.B., Bengio, Y., Liang, J., 2019. Learning fixed points in generative adversarial networks: From image-to-image translation to disease detection and localization. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 191–200.
- Spitzer, H., Kowitz, K., Amunts, K., Harmeling, S., Dickscheid, T., 2018. Improving cytoarchitectonic segmentation of human brain areas with self-supervised siamese networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 663–671.
- Standley, T., Zamir, A.R., Chen, D., Guibas, L., Malik, J., Savarese, S., 2019. Which tasks should be learned together in multi-task learning? arXiv preprint arXiv:1905.07553.
- Sun, C., Guo, S., Zhang, H., Li, J., Chen, M., Ma, S., Jin, L., Liu, X., Li, X., Qian, X., 2017. Automatic segmentation of liver tumors from multiphase contrast-enhanced ct images based on fcns. *Artif. Intell. Med.* 83, 58–66.
- Sun, C., Shrivastava, A., Singh, S., Gupta, A., 2017. Revisiting unreasonable effectiveness of data in deep learning era. In: Proceedings of the IEEE international conference on computer vision, pp. 843–852.
- Sun, W., Zheng, B., Qian, W., 2017. Automatic feature learning using multichannel roi based on deep structured algorithms for computerized lung cancer diagnosis. *Comput. Biol. Med.* 89, 530–539.
- Tajbakhsh, N., Gotway, M.B., Liang, J., 2015. Computer-aided pulmonary embolism detection using a novel vessel-aligned multi-planar image representation and convolutional neural networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 62–69.
- Tajbakhsh, N., Hu, Y., Cao, J., Yan, X., Xiao, Y., Lu, Y., Liang, J., Terzopoulos, D., Ding, X., 2019. Surrogate supervision for medical image analysis: effective deep learning from limited quantities of labeled data. arXiv preprint arXiv:1901.08707.
- Tajbakhsh, N., Shin, J.Y., Gotway, M.B., Liang, J., 2019. Computer-aided detection and visualization of pulmonary embolism using a novel, compact, and discriminative image representation. *Med. Image Anal.* 58, 101541.
- Tajbakhsh, N., Shin, J.Y., Gurudu, S.R., Hurst, R.T., Kendall, C.B., Gotway, M.B., Liang, J., 2016. Convolutional neural networks for medical image analysis: full training or fine tuning? *IEEE Trans. Med. Imaging* 35 (5), 1299–1312.
- Taleb, A., Loetzsch, W., Danz, N., Severin, J., Gaertner, T., Bergner, B., Lipfert, C., 2020. 3D self-supervised methods for medical imaging. arXiv preprint arXiv:2006.03829.
- Tang, H., Zhang, C., Xie, X., 2019. Nodulenet: Decoupled false positive reduction for pulmonary nodule detection and segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 266–274.
- Tang, Y., Wang, X., Harrison, A.P., Lu, L., Xiao, J., Summers, R.M., 2018. Attention-guided curriculum learning for weakly supervised classification and localization of thoracic diseases on chest radiographs. In: International Workshop on Machine Learning in Medical Imaging. Springer, pp. 249–258.
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., Manzagol, P.-A., 2010. Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research* 11 (Dec), 3371–3408.
- Voulodimos, A., Doulamis, N., Doulamis, A., Protopapadakis, E., 2018. Deep learning for computer vision: a brief review. *Comput. Intell. Neurosci.* 2018.
- Wang, H., Zhou, Z., Li, Y., Chen, Z., Lu, P., Wang, W., Liu, W., Yu, L., 2017. Comparison of machine learning methods for classifying mediastinal lymph node metastasis of non-small cell lung cancer from 18 f-dg pet/ct images. *EJNMMI Res.* 7 (1), 11.
- Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M., 2017. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2097–2106.
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., 2004. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* 13 (4), 600–612.
- Weiss, K., Khoshgoftaar, T.M., Wang, D., 2016. A survey of transfer learning. *J. Big Data* 3 (1), 9.
- Wong, S.C., Gatt, A., Stamatescu, V., McDonnell, M.D., 2016. Understanding data augmentation for classification: when to warp? In: 2016 international conference on digital image computing: techniques and applications (DICTA). IEEE, pp. 1–6.
- Wu, B., Zhou, Z., Wang, J., Wang, Y., 2018. Joint learning for pulmonary nodule segmentation, attributes and malignancy prediction. In: 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). IEEE, pp. 1109–1113.
- Wu, S., Zhang, H.R., Valiant, G., Ré, C., 2020. On the generalization effects of linear transformations in data augmentation. arXiv preprint arXiv:2005.00695.
- Yosinski, J., Clune, J., Bengio, Y., Lipson, H., 2014. How transferable are features in deep neural networks? In: Advances in neural information processing systems, pp. 3320–3328.
- Zhang, L., Qi, G.-J., Wang, L., Luo, J., 2019. Aet vs. aed: Unsupervised representation learning by auto-encoding transformations rather than data. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2547–2555.
- Zhang, R., Isola, P., Efros, A.A., 2016. Colorful image colorization. In: Proceedings of the European Conference on Computer Vision. Springer, pp. 649–666.
- Zhang, R., Isola, P., Efros, A.A., 2017. Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1058–1067.
- Zhang, T., 2004. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In: Proceedings of the twenty-first international conference on Machine learning. ACM, p. 116.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A., 2016. Learning deep features for discriminative localization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2921–2929.
- Zhou, Z., Shin, J., Feng, R., Hurst, R.T., Kendall, C.B., Liang, J., 2019. Integrating active learning and transfer learning for carotid intima-media thickness video interpretation. *J. Digit. Imaging* 32 (2), 290–299.
- Zhou, Z., Shin, J., Zhang, L., Gurudu, S., Gotway, M., Liang, J., 2017. Fine-tuning convolutional neural networks for biomedical image analysis: actively and incrementally. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 7340–7349.
- Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J., 2018. Unet++: A Nested U-net Architecture for Medical Image Segmentation. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. Springer, pp. 3–11.
- Zhou, Z., Sodha, V., Rahman Siddiquee, M.M., Feng, R., Tajbakhsh, N., Gotway, M.B., Liang, J., 2019. Models genesis: Generic autodidactic models for 3d medical image analysis. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2019. Springer International Publishing, Cham, pp. 384–393.
- Zhu, J., Li, Y., Hu, Y., Ma, K., Zhou, S.K., Zheng, Y., 2020. Rubiks cube+: a self-supervised feature learning framework for 3d medical image analysis. *Med. Image Anal.* 101746.
- Zhuang, X., Li, Y., Hu, Y., Ma, K., Yang, Y., Zheng, Y., 2019. Self-supervised feature learning for 3d medical images by playing a rubiks cube. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer, pp. 420–428.