Self-Supervised Pre-Training of Swin Transformers for 3D Medical Image Analysis

Yucheng Tang*
Vanderbilt University

Dong Yang
NVIDIA

Wenqi Li
NVIDIA

Holger R. Roth
NVIDIA

Bennett Landman
Vanderbilt University

Daguang Xu
NVIDIA

Vishwesh Nath[†]
NVIDIA

Ali Hatamizadeh[†‡]
NVIDIA

Abstract

Vision Transformers (ViT)s have shown great performance in self-supervised learning of global and local representations that can be transferred to downstream applications. Inspired by these results, we introduce a novel self-supervised learning framework with tailored proxy tasks for medical image analysis. Specifically, we propose: (i) a new 3D transformer-based model, dubbed Swin UNEt TRansformers (Swin UNETR), with a hierarchical encoder for self-supervised pre-training; (ii) tailored proxy tasks for learning the underlying pattern of human anatomy. We demonstrate successful pre-training of the proposed model on 5,050 publicly available computed tomography (CT) images from various body organs. The effectiveness of our approach is validated by fine-tuning the pre-trained models on the Beyond the Cranial Vault (BTCV) Segmentation Challenge with 13 abdominal organs and segmentation tasks from the Medical Segmentation Decathlon (MSD) dataset. Our model is currently the state-of-the-art on the public test leaderboards of both MSD[1] and BTCV[2] datasets. Code: https://monai.io/research/swin-unetr.

1. Introduction

Vision Transformers (ViT)s [16] have started a revolutionary trend in computer vision [12, 54] and medical image analysis [6, 19]. Transformers demonstrate exceptional capability in learning pre-text tasks, are effective in learning of global and local information across layers, and provide

---

*Work completed during internship at NVIDIA

[†]Co-senior advising

[‡]Corresponding author: ahatamizadeh@nvidia.com

[1]https://decathlon-10.grand-challenge.org/evaluation/challenge/leaderboard/

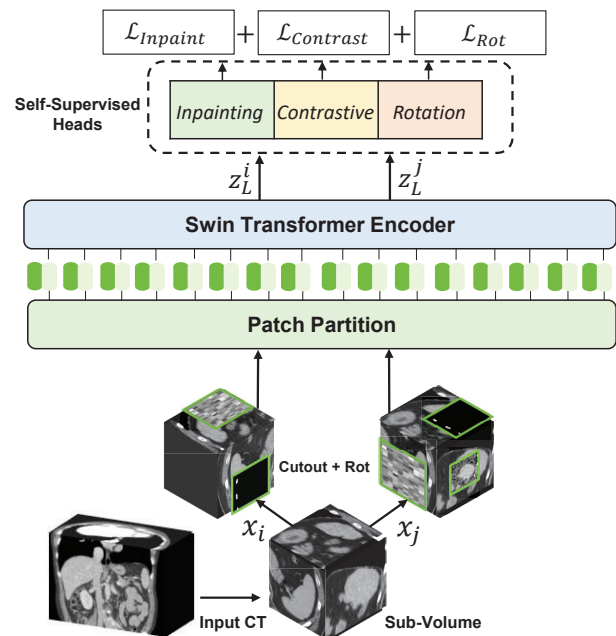[2]https://www.synapse.org/#!Synapse:syn3193805/wiki/217785/

Figure 1. Overview of our proposed pre-training framework. Input CT images are randomly cropped into sub-volumes and augmented with random inner cutout and rotation, then fed to the Swin UNETR encoder as input. We use masked volume inpainting, contrastive learning and rotation prediction as proxy tasks for learning contextual representations of input images.

scalability for large-scale training [38, 52]. As opposed to convolutional neural networks (CNNs) with limited receptive fields, ViTs encode visual representations from a sequence of patches and leverage self-attention blocks for modeling long-range global information [38]. Recently, Shifted windows (Swin) Transformers [30] proposed a hierarchical ViT that allows for local computing of self-attention with

non-overlapping windows. This architecture achieves linear complexity as opposed to quadratic complexity of self-attention layers in ViT, hence making it more efficient. In addition, due to the hierarchical nature of Swin Transformers, they are well-suited for tasks requiring multi-scale modeling.

In comparison to CNN-based counterparts, transformer-based models learn stronger features representations during pre-training, and as a result perform favorably on fine-tuning downstream tasks [38]. Several recent efforts on ViTs [5, 48] have achieved new state-of-the-art results by self-supervised pre-training on large-scale datasets such as ImageNet [15].

In addition, medical image analysis has not benefited from these advances in general computer vision due to: (1) large domain gap between natural images and medical imaging modalities, like computed tomography (CT) and magnetic resonance imaging (MRI); (2) lack of cross-plane contextual information when applied to volumetric (3D) images (such as CT or MRI). The latter is a limitation of 2D transformer models for various medical imaging tasks such as segmentation. Prior studies have demonstrated the effectiveness of supervised pre-training in medical imaging for different applications [9, 39]. But creating expert-annotated 3D medical datasets at scale is a non-trivial and time-consuming effort.

To tackle these limitations, we propose a novel self-supervised learning framework for 3D medical image analysis. First, we propose a new architecture dubbed Swin UNEt TRansformers (Swin UNETR) with a Swin Transformer encoder that directly utilizes 3D input patches. Subsequently, the transformer encoder is pre-trained with tailored, self-supervised tasks by leveraging various proxy tasks such as image inpainting, 3D rotation prediction, and contrastive learning (See Fig. 1 for an overview). Specifically, the human body presents naturally consistent contextual information in radiographic images such as CT due to its depicted anatomical structure [43, 50]. Hence, proxy tasks are utilized for learning the underlying patterns of the human anatomy. For this purpose, we extracted numerous patch queries from different body compositions such as head, neck, lung, abdomen, and pelvis to learn robust feature representations from various anatomical contexts, organs, tissues, and shapes.

Our framework utilizes contrastive learning [35], masked volume inpainting [37], and 3D rotation prediction [17] as pre-training proxy tasks. The contrastive learning is used to differentiate various ROIs of different body compositions, whereas the inpainting allows for learning the texture, structure and correspondence of masked regions to their surrounding context. The rotation task serves as a mechanism to learn the structural content of images and generates various sub-volumes that can be used for contrastive learning. We utilize these proxy tasks to pre-train our proposed framework on a collection of 5,050 CT images that are acquired from various publicly available datasets.

Furthermore, to validate the effectiveness of pre-training,

we use 3D medical image segmentation as a downstream application and reformulate it as a 1D sequence-to-sequence prediction task. For this purpose, we leverage the Swin UNETR encoder with hierarchical feature encoding and shifted windows to extract feature representations at four different resolutions. The extracted representations are then connected to a CNN-based decoder. A segmentation head is attached at the end of the decoder for computing the final segmentation output. We fine-tune Swin UNETR with pre-trained weights on two publicly available benchmarks of Medical Segmentation Decathlon (MSD) and the Beyond the Cranial Vault (BTCV). Our model is currently the state-of-the-art on their respective public test leaderboards.

Our main contributions in this work are summarized as follows:

- We introduce a novel self-supervised learning framework with tailored proxy tasks for pre-training on CT image datasets. To this end, we propose a novel 3D transformer-based architecture, dubbed as Swin UNETR, consisting of an encoder that extracts feature representations at multiple resolutions and is utilized for pre-training.

- We demonstrate successful pre-training on a cohort of 5,050 publicly available CT images from various applications using the proposed encoder and proxy tasks. This results in a powerful pre-trained model with robust feature representation that could be utilized for various medical image analysis downstream tasks.

- We validate the effectiveness of proposed framework by fine-tuning the pre-trained Swin UNETR on two public benchmarks of MSD and BTCV and achieve ***state-of-the-art*** on the test leaderboards of both datasets.

2. Related Works

Medical Segmentation with Transformers Vision transformers are first used in classification tasks and are adopted from sequence-to-sequence modeling in natural language processing. Self-attention mechanisms that aggregate information from the entire input sequence are first achieving comparable, then better performance against prior arts of convolutional architectures such as ResNet [21] or U-Net [13]. Recently, transformer-based networks [24, 47, 49, 55] are proposed for medical image segmentation. In these pioneering works, the transformer blocks are used either as a bottleneck feature encoder or as additional modules after convolutional layers, resulting in limited exploitation of the spatial context advantages of transformers. Comparing to prior works [6, 47], which are using transformers as secondary encoder, we propose to utilize transformers to embed high-dimensional volumetric medical images, which allow for a more direct encoding of 3D patches and positional embeddings.
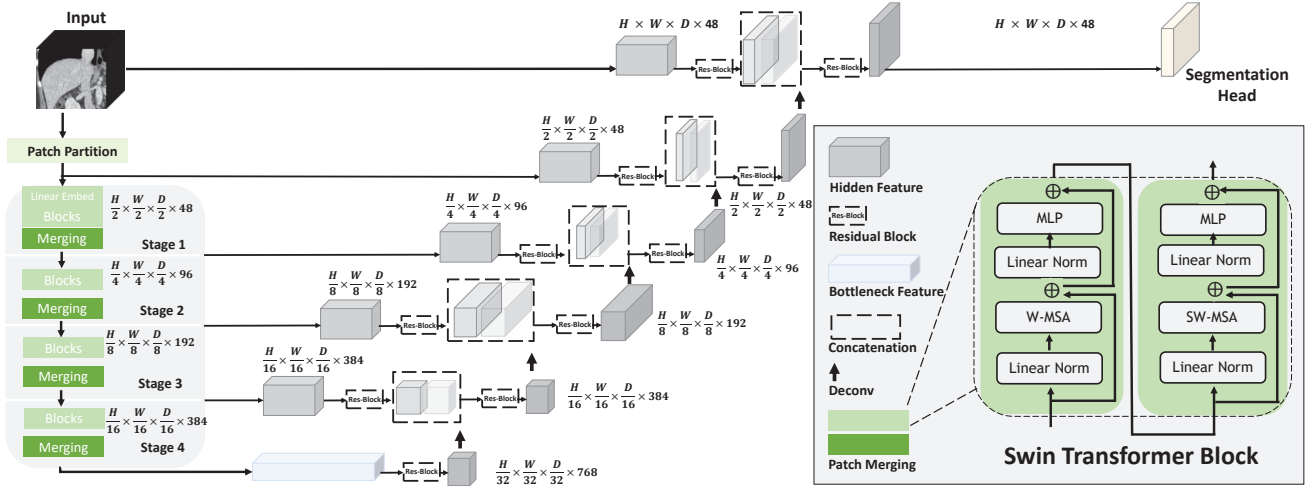
Figure 2. Overview of the Swin UNETR architecture.

Most medical image analysis tasks such as segmentation requires dense inference from multi-scale features. Skip connection-based architectures such as UNet [13] and pyramid networks [40] are widely adopted to leverage hierarchical features. However, vision transformers with a single patch size, while successful in natural image applications, are intractable for high-resolution and high-dimensional volumetric images. To avoid quadratic overflow of computing self-attention at scales [29, 41], Swin Transformer [30, 31] is proposed to construct hierarchical encoding by a shifted-window mechanism. Recent works such as Swin UNet [4] and DS-TransUNet [28] utilize the merits of Swin Transformers for 2D segmentation and achieve promising performance. Augmenting the above-mentioned methods, we learn from 3D anatomy in broader medical image segmentation scenarios by incorporating hierarchically volumetric context.

Pre-training in Medical Image Analysis In medical image analysis, previous studies of pre-training on labeled data demonstrate improved performance by transfer learning [9, 39]. However, generating annotation for medical images is expensive and time-consuming. Recent advances in self-supervised learning offer the promise of utilizing unlabeled data. Self-supervised representation learning [2, 14, 27] constructs feature embedding spaces by designing pre-text tasks, such as solving jigsaw puzzles [34]. Another commonly used pre-text task is to memorize spatial context from medical images, which is motivated by image restoration. This idea is generalized to inpainting tasks [18, 37, 57] to learn visual representations [3, 7, 46] by predicting the original image patches. Similar efforts for reconstructing spatial context have been formulated as solving Rubik's cube problem [58], random rotation prediction [17, 42] and contrastive coding [11, 35]. Different from these efforts, our pre-training framework is simulta-

neously trained with a combination of pre-text tasks, tailored for 3D medical imaging data, and leverages a transformer-based encoder as a powerful feature extractor.

3. Swin UNETR

Swin UNETR comprises a Swin Transformer [30] encoder that directly utilizes 3D patches and is connected to a CNN-based decoder via skip connections at different resolutions. Fig. 2 illustrates the overall architecture of Swin UNETR. We describe the details of encoder and decoder in this section.

3.1. Swin Transformer Encoder

Assuming that the input to the encoder is a sub-volume $\mathcal{X} \in \mathbb{R}^{H \times W \times D \times S}$, a 3D token with a patch resolution of (H', W', D') has a dimension of $H' \times W' \times D' \times S$. The patch partitioning layer creates a sequence of 3D tokens with size $\frac{H}{H'} \times \frac{W}{W'} \times \frac{D}{D'}$ that are projected into a C-dimensional space via an embedding layer. Following [30], for efficient modeling of token interactions, we partition the input volumes into non-overlapping windows and compute local self-attention within each region. Specifically, at layer l, we use a window of size $M \times M \times M$ to evenly divide a 3D token into $\left\lceil \frac{H'}{M} \right\rceil \times \left\lceil \frac{W'}{M} \right\rceil \times \left\lceil \frac{D'}{M} \right\rceil$ windows. In the subsequent layer $l+1$, we shift the partitioned windows by $\left(\left\lfloor \frac{M}{2} \right\rfloor, \left\lfloor \frac{M}{2} \right\rfloor, \left\lfloor \frac{M}{2} \right\rfloor \right)$ voxels. The shifted windowing mechanism is illustrated in Fig. 3. The outputs of encoder blocks in layers l and $l+1$ are computed as in

$$
\begin{aligned}
\hat{z}^l &= \text{W-MSA}(\text{LN}(z^{l-1})) + z^{l-1} \\
z^l &= \text{MLP}(\text{LN}(\hat{z}^l)) + \hat{z}^l \\
\hat{z}^{l+1} &= \text{SW-MSA}(\text{LN}(z^l)) + z^l \\
z^{l+1} &= \text{MLP}(\text{LN}(\hat{z}^{l+1})) + \hat{z}^{l+1},
\end{aligned}
\tag{1}
$$

where W-MSA and SW-MSA denote regular and window partitioning multi-head self-attention modules, respectively,

$\hat{z}^l$ and $\hat{z}^l$ are the outputs of W-MSA and SW-MSA; LN and MLP denote layer normalization and Multi-Layer Perceptron (see Fig. 2). Following [30], we adopt a 3D cyclic-shifting for efficient batch computation of shifted windowing. Furthermore, we calculate the self-attention according to

$$\text{Attention}(Q,K,V) = \text{Softmax}\left(\frac{QK^\top}{\sqrt{d}}\right)V. \qquad (2)$$

where Q,K,V represent queries, keys and values respectively, d is the size of the query and key.

Our encoder uses a patch size of $2 \times 2 \times 2$ with a feature dimension of $2 \times 2 \times 2 \times 1 = 8$ (*i.e.* single input channel CT images) and a $C = 48$-dimensional embedding space. Furthermore, the overall architecture of the encoder consists of 4 stages comprising of 2 transformer blocks at each stage (*i.e.* $L = 8$ total layers). In between every stage, a patch merging layer is used to reduce the resolution by a factor of 2. Stage 1 consists of a linear embedding layer and transformer blocks that maintain the number of tokens as $\frac{H}{2} \times \frac{W}{2} \times \frac{D}{2}$. Furthermore, a patch merging layer groups patches with resolution $2 \times 2 \times 2$ and concatenates them, resulting in a $4C$-dimensional feature embedding. A linear layer is then used to downsample the resolution by reducing the dimension to $2C$. The same procedure continues in stage 2, stage 3 and stage 4 with resolutions of $\frac{H}{4} \times \frac{W}{4} \times \frac{D}{4}$, $\frac{H}{8} \times \frac{W}{8} \times \frac{D}{8}$ and $\frac{H}{16} \times \frac{W}{16} \times \frac{D}{16}$ respectively. The hierarchical representations of the encoder at different stages are used in downstream applications such as segmentation for multi-scale feature extraction.

3.2. Decoder

The encoder of Swin UNETR is connected to a CNN-based decoder at each resolution via skip connections to create a "U-shaped" network for downstream applications such as segmentation. Specifically, we extract the output sequence representations of each stage i $(i \in \{0,1,2,3,4\})$ in the encoder as well as the bottleneck $(i = 5)$ and reshape them into features with size $\frac{H}{2^i} \times \frac{W}{2^i} \times \frac{D}{2^i}$. The extracted representations at each stage are then fed into a residual block consisting of two post-normalized $3 \times 3 \times 3$ convolutional layers with instance normalization [45]. The processed features from each stage are then upsampled by using a deconvolutional layer and concatenated with processed features of the preceding stage. The concatenated features are fed into a residual block with aforementioned descriptions. For segmentation, we concatenate the output of the encoder (*i.e.* Swin Transformer) with processed features of the input volume and feed them into a residual block followed by a final $1 \times 1 \times 1$ convolutional layer with a proper activation function (*i.e.* softmax) for computing the segmentation probabilities (see Fig. 2 for details of the architecture).

4. Pre-training

We pre-train the Swin UNETR encoder with multiple proxy tasks and formulate it with a multi-objective loss func-

3D Tokens: 8× 8 × 8
Window size: 4× 4 × 4
Layer *l*
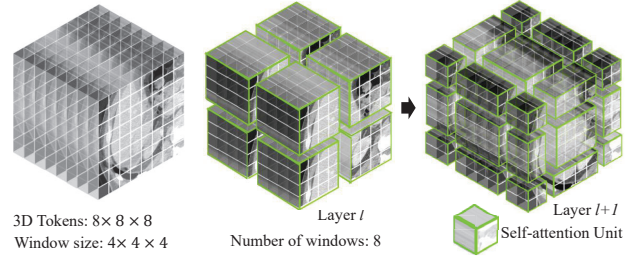Number of windows: 8
Layer *l+1*
Self-attention Unit

Figure 3. Shifted windowing mechanism for efficient self-attention computation of 3D tokens with $8 \times 8 \times 8$ tokens and $4 \times 4 \times 4$ window size.

tion (Fig. 1). The objective of self-supervised representation learning is to encode region of interests (ROI)-aware information of the human body. Inspired by previous works on context reconstruction [18,57] and contrastive encoding [20], we exploit three proxy tasks for medical image representation learning. Three additional projection heads are attached to the encoder during pre-training. Furthermore, the downstream task, e.g. segmentation, fine-tunes the full Swin UNETR model with the projection heads removed. In training, sub-volumes are cropped random regions of the volumetric data. Then, stochastic data augmentations with random rotation and cutout are applied twice to each sub-volume within a mini-batch, resulting in two views of each data.

4.1. Masked Volume Inpainting

The cutout augmentation masks out ROIs in the sub-volume $\mathcal{X} \in \mathbb{R}^{H \times W \times D \times C}$ randomly with volume ratio of s. We attach a transpose convolution layer to the encoder as the reconstruction head and denote its output as $\hat{\mathcal{X}}^\mathcal{M}$. The reconstruction objective is defined by an $L1$ loss between $\mathcal{X}$ and $\hat{\mathcal{X}}^\mathcal{M}$

$$\mathcal{L}_{inpaint} = \|\mathcal{X} - \hat{\mathcal{X}}^\mathcal{M}\|_1, \qquad (3)$$

The masked volume inpainting is motivated by prior work which focused on 2D images [37]. We extend it to 3D domain to showcase its effectiveness on representation learning of volumetric medical images.

4.2. Image Rotation

The rotation prediction task predicts the angle categories by which the input sub-volume is rotated. For simplicity, we employ R classes of $0°, 90°, 180°, 270°$ rotations along the z-axis. An MLP classification head is used for predicting the softmax probabilities $\hat{y}_r$ of rotation categories. Given the ground truth y_r, a cross-entropy loss is used for rotation prediction task:

$$\mathcal{L}_{rot} = -\sum_{r=1}^{R} y_r \log(\hat{y}_r), \qquad (4)$$

The 3D rotation and cutout also serves simultaneously as an augmentation transformation for contrastive learning.

4.3. Contrastive Coding

The self-supervised contrastive coding presents promising performance on visual representation learning when transferred to downstream tasks [10, 36]. Given a batch of augmented sub-volumes, the contrastive coding allows for a better representation learning by maximizing the mutual information between positive pairs (augmented samples from same sub-volume), while minimizing that between negative pairs (views from different sub-volumes). The contrastive coding is obtained by attaching a linear layer to the Swin UNETR encoder, which maps each augmented sub-volume to a latent representation v. We use cosine similarity as the distance measurement of the encoded representations as defined in [10]. Formally, the 3D contrastive coding loss between a pair v_i and v_j is defined as:

$$\mathcal{L}_{contrast} = -\log \frac{\exp(sim(v_i, v_j)/t)}{\sum_k^{2N} 1_{k \neq i} \exp(sim(v_i, v_k)/t)}, \quad (5)$$

where t is the measurement of normalized temperature scale. 1 is the indicator function evaluating to 1 iff $k \neq i$. sim denotes the dot product between normalized embeddings. The contrastive learning loss function strengthens the intra-class compactness as well as the inter-class separability.

4.4. Loss Function

Formally, we minimize the total loss function by training Swin UNETR's encoder with multiple pre-training objectives of masked volume inpainting, 3D image rotation & contrastive coding as follows:

$$\mathcal{L}_{tot} = \lambda_1 \mathcal{L}_{inpaint} + \lambda_2 \mathcal{L}_{contrast} + \lambda_3 \mathcal{L}_{rot}. \quad (6)$$

A grid-search hyper-parameter optimization was performed which estimated the optimal values of $\lambda_1 = \lambda_2 = \lambda_3 = 1$.

5. Experiments

5.1. Datasets

Pre-training Datasets : A total of 5 public CT datasets, consisting of 5,050 subjects, are used to construct our pre-training dataset. The corresponding number of 3D volumes for chest, abdomen and head/neck are 2,018, 1,520 and 1,223 respectively. The collection and source details are presented in the supplementary materials. Existing annotations or labels are *not* utilized from these datasets during the pre-training stage.

BTCV : The Beyond the Cranial Vault (BTCV) abdomen challenge dataset [26] contains 30 subjects with abdominal CT scans where 13 organs are annotated by interpreters under supervision of radiologists at Vanderbilt University Medical Center. Each CT scan is acquired with contrast enhancement phase at portal venous consists of 80 to 225 slices with 512×512 pixels and slice thickness ranging from 1 to 6 mm. The multi-organ segmentation problem is formulated as a 13 classes segmentation task (see Table 1 for details). The pre-processing pipeline is detailed in supplementary materials.

MSD: Medical Segmentation Decathlon (MSD) dataset [1] comprises of 10 segmentation tasks from different organs and image modalities. These tasks are designed to feature difficulties across medical images, such as small training sets, unbalanced classes, multi-modality data and small objects. Therefore, the MSD challenge can serve as a comprehensive benchmark to evaluate the generalizability of medical image segmentation methods. The pre-processing pipeline for this dataset is outlined in supplementary materials.

5.2. Implementation Details

For pre-training tasks, (1) masked volume inpainting: the ROI dropping rate is set to 30% (as also used in [2]); the dropped regions are randomly generated and they sum up to reach overall number of voxels; (2) 3D contrastive coding: a feature size of 512 is used as the embedding size; (3) rotation prediction: the rotation degree is configured to $0°$, $90°$, $180°$, and $270°$. We train the model using the AdamW [32] optimizer with a warm-up cosine scheduler of 500 iterations. The pre-training experiments use a batch-size of 4 per GPU (with $96 \times 96 \times 96$ patch), and initial learning rate of $4e^{-4}$, momentum of 0.9 and decay of $1e^{-5}$ for 450K iterations. Our model is implemented in PyTorch and MONAI[4]. A five-fold cross validation strategy is used to train models for all BTCV and MSD experiments. We select the best model in each fold and ensemble their outputs for final segmentation predictions. Detailed training hyperparameters for fine-tuning BTCV and MSD tasks can be found in the supplementary materials. All models are trained on a NVIDIA DGX-1 server.

5.3. Evaluation Metrics

The Dice similarity coefficient (Dice) and Hausdorff Distance 95% (HD95) are used as measurements for experiment results. HD95 calculates 95^{th} percentile of surface distances between ground truth and prediction point sets. Metric formulations are as follows:

$$\text{Dice} = \frac{2 \sum_{i=1}^I Y_i \hat{Y}_i}{\sum_{i=1}^I Y_i + \sum_{i=1}^I \hat{Y}_i}, \quad (7)$$

$$\text{HD} = \max\{\max_{y' \in Y'} \min_{\bar{y}' \in \bar{Y}'} \|y' - \bar{y}'\|, \max_{\bar{y}' \in \bar{Y}'} \min_{y' \in Y'} \|\bar{y}' - y'\|\}. \quad (8)$$

where Y and $\bar{Y}$ denote the ground truth and prediction of voxel values. Y' and $\hat{Y}'$ denote ground truth and prediction

---

[3] https://www.synapse.org/#!Synapse:syn3193805/wiki/217785/

[4] https://monai.io/

Methods	Spl	RKid	LKid	Gall	Eso	Liv	Sto	Aor	IVC	Veins	Pan	AG	Avg.
SETR NUP [53]	0.931	0.890	0.897	0.652	0.760	0.952	0.809	0.867	0.745	0.717	0.719	0.620	0.796
SETR PUP [53]	0.929	0.893	0.892	0.649	0.764	0.954	0.822	0.869	0.742	0.715	0.714	0.618	0.797
SETR MLA [54]	0.930	0.889	0.894	0.650	0.762	0.953	0.819	0.872	0.739	0.720	0.716	0.614	0.796
ASPP [8]	0.935	0.892	0.914	0.689	0.760	0.953	0.812	0.918	0.807	0.695	0.720	0.629	0.811
TransUNet [6]	0.952	0.927	0.929	0.662	0.757	0.969	0.889	0.920	0.833	0.791	0.775	0.637	0.838
CoTr* [47]	0.943	0.924	0.929	0.687	0.762	0.962	0.894	0.914	0.838	0.796	0.783	0.647	0.841
CoTr [47]	0.958	0.921	0.936	0.700	0.764	0.963	0.854	0.920	0.838	0.787	0.775	0.694	0.844
RandomPatch [44]	0.963	0.912	0.921	0.749	0.760	0.962	0.870	0.889	0.846	0.786	0.762	0.712	0.844
PaNN [56]	0.966	0.927	0.952	0.732	0.791	0.973	0.891	0.914	0.850	0.805	0.802	0.652	0.854
nnUNet [23]	0.967	0.924	**0.957**	0.814	0.832	0.975	0.925	0.928	0.870	0.832	0.849	0.784	0.888
UNETR [19]	0.972	0.942	0.954	0.825	0.864	0.983	0.945	**0.948**	0.890	0.858	0.852	0.812	0.891
Swin UNETR	**0.976**	**0.958**	0.956	**0.893**	**0.875**	**0.985**	**0.953**	**0.948**	**0.904**	**0.899**	**0.897**	**0.846**	**0.918**

Table 1. Leaderboard[3]Dice results of BTCV challenge on multi-organ segmentation. The proposed method achieves state-of-the-art performance in both free and standard competitions. Note: Spl: spleen, RKid: right kidney, LKid: left kidney, Gall: gallbladder, Eso: esophagus, Liv: liver, Sto: stomach, Aor: aorta, IVC: inferior vena cava, Veins: portal and splenic veins, Pan: pancreas, AG: left and right adrenal glands.
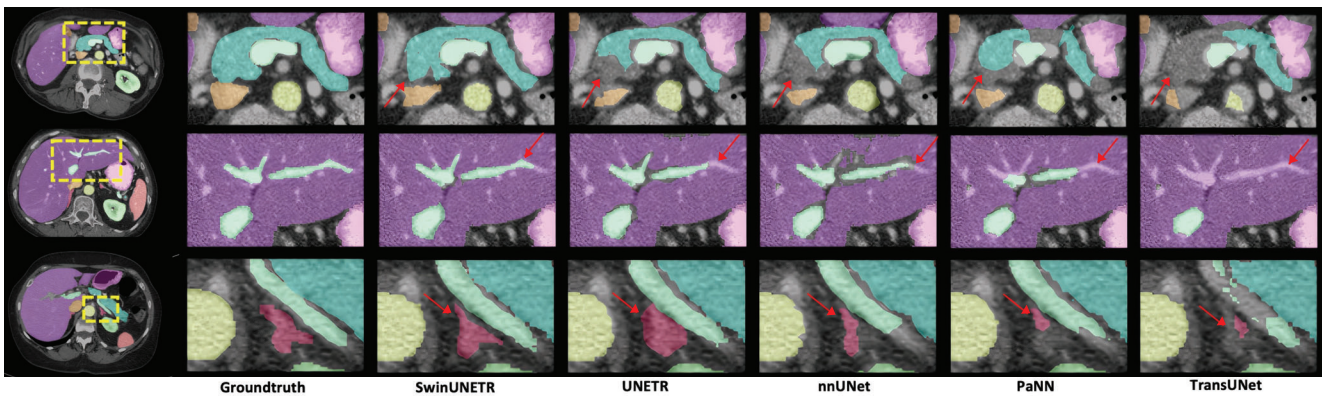

Figure 4. Qualitative visualizations of the proposed Swin UNETR and baseline methods. Three representative subjects are demonstrated. Regions of evident improvements are enlarged to show better details of pancreas (blue), portal vein (light green), and adrenal gland (red).

Method	Rank	Average Accuracy	
		Dice ↑	NSD ↑
Swin UNETR	1	**78.68**	**89.28**
DiNTS [22]	2	77.93	88.68
nnUNet [23]	3	77.89	88.09
Models Gen. [57]	4	76.97	87.19
Trans VW [18]	5	76.96	87.64

Table 2. Overall performance of top-ranking methods on all 10 segmentation tasks in the MSD public test leaderboard. NSD denotes Normalized Surface Distance.

surface point sets. Surface Dice [33] is also used, which is referred as Normalized Surface Distance (NSD) in MSD challenge evaluation. The metric measures the overlap of ground truth and prediction surfaces (with a fixed tolerance) instead of the overlap of two volumes. This provides a measure of agreement between the surfaces of two structures.

5.4. Results

5.4.1 BTCV Multi-organ Segmentation Challenge

We extensively compare the benchmarks of our model with baselines. The published leaderboard evaluation is shown in Table 1. Compared with other top submissions, the proposed Swin UNETR achieves the best performance. We obtain the state-of-the-art Dice of 0.908, outperforming the second, third and fourth top-ranked baselines by 1.6%, 2.0% and 2.4% on average of 13 organs, respectively. Distinct improvements can be specifically observed for organs that are smaller in size, such as splenic and portal veins of 3.6% against prior state-of-the-art method, pancreas of 1.6%, and adrenal glands of 3.8%. Moderate improvements are observed in other organs. The representative samples in Fig. 4 demonstrate the success of identifying organ details by Swin UNETR. Our method detects the pancreas tail (row 1), and branches in the portal vein (row 2) in Fig. 4, where other methods under segment parts of each tissue. In addition, our method demonstrates distinct improvement in segmentation of adrenal glands (row 3).

Organ	Task01 Brain Tumour								Task03 Liver						Task06 Lung	
Metric	Dice1	Dice2	Dice3	Avg.	NSD1	NSD2	NSD3	Avg.	Dice1	Dice2	Avg.	NSD1	NSD2	Avg.	Dice1	NSD1
Kim et al [25]	67.40	45.75	68.26	60.47	86.65	72.03	90.28	82.99	94.25	72.96	83.61	96.76	88.58	92.67	63.10	62.51
Trans VW [18]	68.03	46.98	68.40	61.14	87.52	72.42	90.91	83.62	95.18	76.90	86.04	97.86	92.03	94.95	74.54	76.22
C2FNAS [51]	67.62	48.60	69.72	61.98	87.61	72.87	91.16	83.88	94.98	72.89	83.94	98.38	89.15	93.77	70.44	72.22
Models Gen. [57]	68.03	46.98	68.40	61.14	87.52	72.42	90.91	83.62	95.72	**77.50**	**86.61**	98.48	91.92	**95.20**	74.54	76.22
nnUNet [23]	68.04	46.81	68.46	61.10	87.51	72.47	90.78	83.59	**95.75**	75.97	85.86	98.55	90.65	94.60	73.97	76.02
DiNTS [22]	69.28	48.65	69.75	62.56	**89.33**	73.16	91.69	84.73	95.35	74.62	84.99	**98.69**	91.02	94.86	74.75	77.02
Swin UNETR	**70.02**	**52.52**	**70.51**	**64.35**	89.07	**80.30**	**93.46**	**87.61**	95.35	75.68	85.52	98.34	91.59	94.97	**76.60**	**77.40**

Organ	Task07 Pancreas						Task08 Hepatic Vessel						Task09 Spleen		Task10 Colon	
Metric	Dice1	Dice2	Avg.	NSD1	NSD2	Avg.	Dice1	Dice2	Avg.	NSD1	NSD2	Avg.	Dice1	NSD1	Dice1	NSD1
Kim et al [25]	80.61	51.75	66.18	95.83	73.09	84.46	62.34	68.63	65.49	83.22	78.43	80.83	91.92	94.83	49.32	62.21
Trans VW [18]	81.42	51.08	66.25	96.07	70.13	83.10	65.80	71.44	68.62	84.01	80.15	82.08	97.35	99.87	51.47	60.53
C2FNAS [51]	80.76	54.41	67.59	96.16	75.58	85.87	64.30	71.00	67.65	83.78	80.66	82.22	96.28	97.66	58.90	**72.56**
Models Gen. [57]	81.36	50.36	65.86	96.16	70.02	83.09	65.80	71.44	68.62	84.01	80.15	82.08	97.35	99.87	51.47	60.53
nnUNet [23]	81.64	52.78	67.21	96.14	71.47	83.81	**66.46**	71.78	**69.12**	84.43	80.72	82.58	**97.43**	**99.89**	58.33	68.43
DiNTS [22]	81.02	55.35	68.19	96.26	75.90	86.08	64.50	71.76	68.13	83.98	81.03	82.51	96.98	99.83	59.21	70.34
Swin UNETR	**81.85**	**58.21**	**70.71**	**96.57**	**79.10**	**87.84**	65.69	**72.20**	68.95	**84.83**	**81.62**	**83.23**	96.99	99.84	**59.45**	70.89

Table 3. MSD test dataset performance comparison of Dice and NSD. Benchmarks obtained from MSD test leaderboard[5].

5.4.2 Segmentation Results on MSD

The overall MSD results per task and ranking from the challenge leaderboard are shown in Table. 2. The proposed Swin UNETR achieves state-of-the-art performance in Task01 BrainTumour, Task06 Lung, Task07 Pancreas, and Task10 Colon. The results are comparable for Task02 Heart, Task03 Liver, Task04 Hippocampus, Task05 Prostate, Task08 HepaticVessel and Task09 Spleen. Overall, Swin UNETR presents the best average Dice of 78.68% across all ten tasks and achieves the top ranking in the MSD leaderboard. The detail number of multiple tasks are shown in Table 3. Qualitative visualization can be observed in Fig. 5. Swin UNETR with self-supervised pre-training demonstrates visually better segmentation results in the CT tasks. The pre-trained weights are only used for fine-tuning CT tasks including Liver, Lung, Pancreas, HepaticVessel, Spleen, and Colon. For MRI tasks: Brain Tumour, Heart, Hippocampus, Prostate, experiments are trained from scratch because of the domain gap between CT and MRI images. Due to space limitations, we present the MSD test benchmarks for the remaining three MRI tasks in the supplementary materials.

5.5. Ablation Study

5.5.1 Efficacy of Pre-training

A comparison of all MSD CT tasks using pre-trained model against training from scratch can be observed in Fig. 6. Distinct improvement can be observed for Task03 Liver, Dice of 77.77% comparing to 75.27%. Task08 Hepatic Vessel achieves 68.52% against 64.63%. Task10 Colon shows the largest improvement, from 34.83% to 43.38%. Task07 Pancreas and Task09 Spleen both achieve significant improvement from 67.12% to 67.82%, and 96.05% to 97.32% respectively.
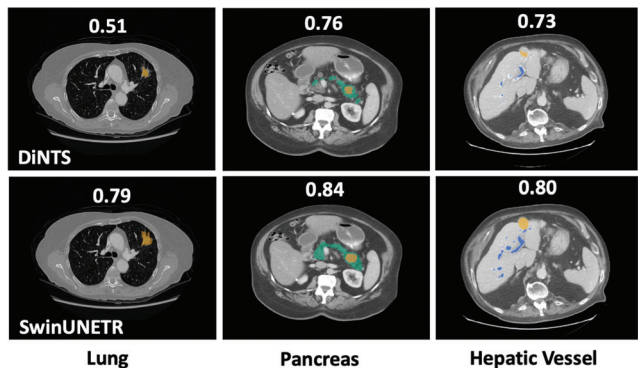

Figure 5. Qualitative results of representative MSD CT tasks. Average Dice values are illustrated on top of each image. Our model demonstrates more accurate performance in comparison to DiNTS for both organ and tumor segmentation across different tasks.

Loss Function	Average Accuracy	
	Dice ↑	HD ↓
Scratch	83.43	42.36
$\mathcal{L}_{rot}$	83.56	36.19
$\mathcal{L}_{contrast}$	83.67	38.81
$\mathcal{L}_{inpaint}$	83.85	28.94
$\mathcal{L}_{inpaint} + \mathcal{L}_{rot}$	84.01	26.06
$\mathcal{L}_{inpaint} + \mathcal{L}_{contrast}$	84.45	24.37
$\mathcal{L}_{inpaint} + \mathcal{L}_{contrast} + \mathcal{L}_{rot}$	**84.72**	**20.03**

Table 4. Ablation study of the effectiveness of each objective function in the proposed pre-training loss. HD denotes Hausdorff Distance. Experiments on fine-tuning the BTCV dataset.

5.5.2 Reduce Manual Labeling Efforts

Fig. 7 demonstrates the comparison results of fine-tuning using a subset of BTCV dataset. We show using 10% of
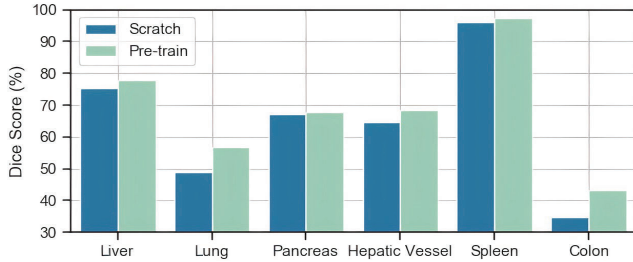
Figure 6. The indication of Dice gap between using pre-training (Green) and scratch model (Blue) on MSD CT tasks validation set.
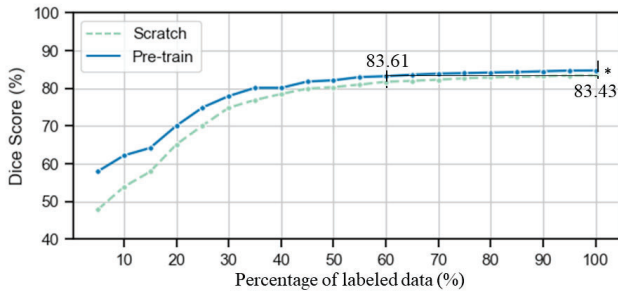

Figure 7. Data-efficient performance on BTCV test dataset. Significance under Wilcoxon Signed Rank test, $*: p < 0.001$.

labeled data, experiments with pre-training weights achieve approximately 10% improvement comparing to training from scratch. On employing all labeled data, the self-supervised pre-training shows 1.3% higher average Dice. The Dice number 83.13 of learning from scratch with entire dataset can be achieved by using pre-trained Swin UNETR with 60% data. Fig. 7 indicates that our approach can reduce the annotation effort by at least 40% for BTCV task.

5.5.3   Size of Pre-training Dataset

We perform organ-wise study on BTCV dataset by using pre-trained weights of smaller unlabeled data. In Fig. 8, the fine-tuning results are obtained from pre-training 100, 3,000, and 5,000 scans. We observe that Swin UNETR is robust with respect to the total number of CT scans trained. Fig. 8 demonstrates the proposed model can benefit from larger pre-training datasets with increasing size of unlabeled data.

5.5.4   Efficacy of Self-Supervised Objectives

We perform empirical study on pre-training with different combinations of self-supervised objectives. As shown in Table 4, on BTCV test set, using pre-trained weights by inpainting achieves the highest improvement at single task modeling. On pairing tasks, inpainting and contrastive learning show Dice of 84.45% and Hausdorff Distance (HD) of 24.37. Overall, employing all proxy tasks achieves best Dice of 84.72%.
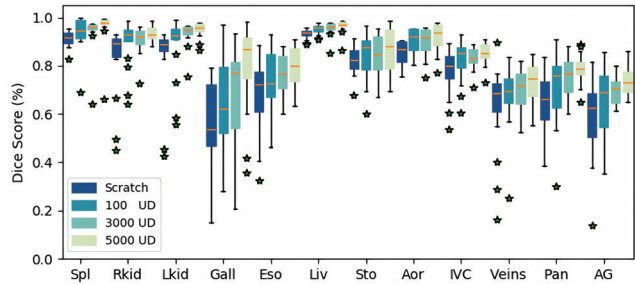

Figure 8. Pre-trained weights using 100, 3000 and 5000 scans are compared for fine-tuning on the BTCV dataset for each organ.

6. Discussion and Limitations

Our state-of-the-art results on the test leaderboards of MSD and BTCV datasets validate the effectiveness of the proposed self-supervised learning framework in taking the advantage of large number of available medical images without the need of annotation effort. Subsequently, fine-tuning the pretrained Swin UNETR model achieves higher accuracy, improves the convergence speed, and reduces the annotation effort in comparison to training with randomly initialized weights from scratch. Our framework is scalable and can be easily extended with more proxy tasks and augmentation transformations. Meanwhile, the pre-trained encoder can benefit the transfer learning of various medical imaging analysis tasks, such as classification and detection. In MSD pancreas segmentation task, Swin UNETR with pre-trained weights outperforms AutoML algorithms such as DiNTS [22] and C2FNAS [51] that are specifically designed for searching the optimal network architectures on the same segmentation task. Currently, Swin UNETR has only been pre-trained using CT images, and our experiments have not demonstrated enough transferability when applied directly to other medical imaging modalities such as MRI. This is mainly due to obvious domain gaps and different number of input channels that are specific to each modality. As a result, this is a potential direction that should be studied in future efforts.

7. Conclusions

In this work, we present a novel framework for self-supervised pre-training of 3D medical images. Inspired by merging feature maps at scales, we built the Swin UNETR by exploiting transformer-encoded spatial representations into convolution-based decoders. By proposing the first transformer-based 3D medical image pre-training, we leverage the power of Swin Transformer encoder for fine-tuning segmentation tasks. Swin UNETR with self-supervised pre-training achieves the state-of-the-art performance on the BTCV multi-organ segmentation challenge and MSD challenge. Particularly, we presented the large-scale CT pre-training with 5,050 volumes, by combining multiple publicly available datasets and diversities of anatomical ROIs.

References

[1] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, Bram van Ginneken, et al. The medical segmentation decathlon. *arXiv preprint arXiv:2106.05735*, 2021. 5

[2] Sara Atito, Muhammad Awais, and Josef Kittler. Sit: Self-supervised vision transformer. *arXiv preprint arXiv:2104.03602*, 2021. 3, 5

[3] Shekoofeh Azizi, Basil Mustafa, Fiona Ryan, Zachary Beaver, Jan Freyberg, Jonathan Deaton, Aaron Loh, Alan Karthikesalingam, Simon Kornblith, Ting Chen, et al. Big self-supervised models advance medical image classification. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 3

[4] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. Swin-unet: Unet-like pure transformer for medical image segmentation. *arXiv preprint arXiv:2105.05537*, 2021. 3

[5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2

[6] Jieneng Chen, Yongyi Lu, Qihang Yu, Xiangde Luo, Ehsan Adeli, Yan Wang, Le Lu, Alan L Yuille, and Yuyin Zhou. Transunet: Transformers make strong encoders for medical image segmentation. *arXiv preprint arXiv:2102.04306*, 2021. 1, 2, 6

[7] Liang Chen, Paul Bentley, Kensaku Mori, Kazunari Misawa, Michitaka Fujiwara, and Daniel Rueckert. Self-supervised learning for medical image analysis using image context restoration. *Medical image analysis*, 58:101539, 2019. 3

[8] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *European Conference on Computer Vision*, 2018. 6

[9] Sihong Chen, Kai Ma, and Yefeng Zheng. Med3d: Transfer learning for 3d medical image analysis. *arXiv preprint arXiv:1904.00625*, 2019. 2, 3

[10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 2020. 5

[11] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 3

[12] Bowen Cheng, Alexander G Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *arXiv preprint arXiv:2107.06278*, 2021. 1

[13] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, 2016. 2, 3

[14] Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. Up-detr: Unsupervised pre-training for object detection with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 3

[15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2009. 2

[16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 1

[17] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018. 2, 3

[18] Fatemeh Haghighi, Mohammad Reza Hosseinzadeh Taher, Zongwei Zhou, Michael B Gotway, and Jianming Liang. Transferable visual words: Exploiting the semantics of anatomical patterns for self-supervised learning. *IEEE transactions on medical imaging*, 2021. 3, 4, 6, 7

[19] Ali Hatamizadeh, Yucheng Tang, Vishwesh Nath, Dong Yang, Andriy Myronenko, Bennett Landman, Holger R Roth, and Daguang Xu. Unetr: Transformers for 3d medical image segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 574–584, 2022. 1, 6

[20] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 4

[21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 2

[22] Yufan He, Dong Yang, Holger Roth, Can Zhao, and Daguang Xu. Dints: Differentiable neural network topology search for 3d medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 6, 7, 8

[23] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2):203–211, 2021. 6, 7

[24] JM Jose and P Oza. Medical transformer: gated axial-attention for medical image segmentation. In *International Conference on Medical Image Computing and Computer Assisted Intervention*, 2021. 2

[25] Sungwoong Kim, Ildoo Kim, Sungbin Lim, Woonhyuk Baek, Chiheon Kim, Hyungjoo Cho, Boogeon Yoon, and Taesup Kim. Scalable neural architecture search for 3d medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2019. 7

[26] B Landman, Z Xu, J Igelsias, M Styner, T Langerak, and A Klein. Miccai multi-atlas labeling beyond the cranial vault–workshop and challenge. In *Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault—Workshop Challenge*, 2015. 5

[27] Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 3

[28] Ailiang Lin, Bingzhi Chen, Jiayu Xu, Zheng Zhang, and Guangming Lu. Ds-transunet: Dual swin transformer u-net for medical image segmentation. *arXiv preprint arXiv:2106.06716*, 2021. 3

[29] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *In Proceedings of the IEEE/CVF Interna-tional Conference on Computer Vision*, 2017. 3

[30] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 1, 3, 4

[31] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. 3

[32] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. 5

[33] Stanislav Nikolov, Sam Blackwell, Alexei Zverovitch, Ruheena Mendes, Michelle Livne, Jeffrey De Fauw, Yojan Patel, Clemens Meyer, Harry Askham, Bernardino Romera-Paredes, et al. Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy. *arXiv preprint arXiv:1809.04430*, 2018. 6

[34] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *European conference on computer vision*, 2016. 3

[35] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 2, 3

[36] Taesung Park, Alexei A Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *European Conference on Computer Vision*, 2020. 5

[37] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016. 2, 3, 4

[38] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *arXiv preprint arXiv:2108.08810*, 2021. 1, 2

[39] Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. *Advances in Neural Information Processing Systems*, 2019. 2, 3

[40] Holger R Roth, Chen Shen, Hirohisa Oda, Takaaki Sugino, Masahiro Oda, Yuichiro Hayashi, Kazunari Misawa, and Kensaku Mori. A multi-scale pyramid of 3d fully convolutional networks for abdominal multi-organ segmentation. In *International conference on medical image computing and computer-assisted intervention*, 2018. 3

[41] Bharat Singh, Mahyar Najibi, and Larry S Davis. Sniper: Efficient multi-scale training. *Advances in Neural Information Processing Systems*, 31:9310–9320, 2018. 3

[42] Aiham Taleb, Winfried Loetzsch, Noel Danz, Julius Severin, Thomas Gaertner, Benjamin Bergner, and Christoph Lippert. 3d self-supervised methods for medical imaging. In *Advances in Neural Information Processing Systems*, 2020. 3

[43] Yucheng Tang, Riqiang Gao, Shizhong Han, Yunqiang Chen, Dashan Gao, Vishwesh Nath, Camilo Bermudez, Michael R Savona, Shunxing Bao, Ilwoo Lyu, et al. Body part regression with self-supervision. *IEEE Transactions on Medical Imaging*, 40(5):1499–1507, 2021. 2

[44] Yucheng Tang, Riqiang Gao, Ho Hin Lee, Shizhong Han, Yunqiang Chen, Dashan Gao, Vishwesh Nath, Camilo Bermudez, Michael R Savona, Richard G Abramson, et al. High-resolution 3d abdominal segmentation with random patch network fusion. *Medical Image Analysis*, 69:101894, 2021. 6

[45] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016. 4

[46] Xiaosong Wang, Ziyue Xu, Leo Tam, Dong Yang, and Daguang Xu. Self-supervised image-text pre-training with mixed data in chest x-rays. *arXiv preprint arXiv:2103.16022*, 2021. 3

[47] Yutong Xie, Jianpeng Zhang, Chunhua Shen, and Yong Xia. Cotr: Efficiently bridging cnn and transformer for 3d medical image segmentation. *International conference on medical image computing and computer-assisted intervention*, 2021. 2, 6

[48] Zhenda Xie, Yutong Lin, Zhuliang Yao, Zheng Zhang, Qi Dai, Yue Cao, and Han Hu. Self-supervised learning with swin transformers. *arXiv preprint arXiv:2105.04553*, 2021. 2

[49] Guoping Xu, Xingrong Wu, Xuan Zhang, and Xinwei He. Levit-unet: Make faster encoders with transformer for medical image segmentation. *arXiv preprint arXiv:2107.08623*, 2021. 2

[50] Ke Yan, Jinzheng Cai, Dakai Jin, Shun Miao, Adam P Harrison, Dazhou Guo, Youbao Tang, Jing Xiao, Jingjing Lu, and Le Lu. Self-supervised learning of pixel-wise anatomical embeddings in radiological images. *arXiv preprint arXiv:2012.02383*, 2020. 2

[51] Qihang Yu, Dong Yang, Holger Roth, Yutong Bai, Yixiao Zhang, Alan L Yuille, and Daguang Xu. C2fnas: Coarse-to-fine neural architecture search for 3d medical image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020. 7, 8

[52] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. *arXiv preprint arXiv:2106.04560*, 2021. 1

[53] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. *arXiv preprint arXiv:2012.15840*, 2020. 6

[54] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers.

In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 1, 6

[55] Hong-Yu Zhou, Jiansen Guo, Yinghao Zhang, Lequan Yu, Liansheng Wang, and Yizhou Yu. nnformer: Interleaved transformer for volumetric segmentation. *arXiv preprint arXiv:2109.03201*, 2021. 2

[56] Yuyin Zhou, Zhe Li, Song Bai, Chong Wang, Xinlei Chen, Mei Han, Elliot Fishman, and Alan L Yuille. Prior-aware neural network for partially-supervised multi-organ segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. 6

[57] Zongwei Zhou, Vatsal Sodha, Jiaxuan Pang, Michael B Gotway, and Jianming Liang. Models genesis. *Medical image analysis*, 67:101840, 2021. 3, 4, 6, 7

[58] Jiuwen Zhu, Yuexiang Li, Yifan Hu, Kai Ma, S Kevin Zhou, and Yefeng Zheng. Rubik's cube+: A self-supervised feature learning framework for 3d medical image analysis. *Medical image analysis*, 64:101746, 2020. 3