

# Which images to label for few-shot medical landmark detection?

Quan Quan<sup>1</sup> Qingsong Yao<sup>1</sup> Jun Li<sup>1</sup> S.Kevin Zhou<sup>2,1</sup>

<sup>1</sup> Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology

<sup>2</sup> University of Science and Technology of China

## Abstract

The success of deep learning methods relies on the availability of well-labeled large-scale datasets. However, for medical images, annotating such abundant training data often requires experienced radiologists and consumes their limited time. Few-shot learning is developed to alleviate this burden, which achieves competitive performances with only several labeled data. However, a crucial yet **previously overlooked** problem in few-shot learning is about the selection of template images for annotation before learning, which affects the final performance. We herein propose a novel Sample Choosing Policy (SCP) to select “the most worthy” images for annotation, in the context of few-shot medical landmark detection. SCP consists of three parts: 1) Self-supervised training for building a pre-trained deep model to extract features from radiological images, 2) Key Point Proposal for localizing informative patches, and 3) Representative Score Estimation for searching the most representative samples or templates. The advantage of SCP is demonstrated by various experiments on three widely-used public datasets. For one-shot medical landmark detection, its use reduces the mean radial errors on Cephalometric and HandXray datasets by 14.2% (from 3.595mm to 3.083mm) and 35.5% (4.114mm to 2.653mm), respectively.

## 1. Introduction

It is widely known that the success of deep learning relies on data availability. Learning from datasets of a larger quantity and higher quality likely brings a higher performance and better generalization for neural networks. Yet, the labeling of datasets needs well-trained, highly-engaged radiologists for medical image analysis tasks [48,49], which is especially challenging as physicians are costly and busy.

To alleviate this problem, many researchers [3,17,19,34] utilize labeled data together with unlabeled data in a semi-supervised style to boost performance. A classic method is mean teacher [17,34], which aggregates multiple predictions of unlabeled data by a teacher model pre-trained from




Figure 1. The distribution of the mean radial error (MRE) when choosing a different image as a template in one-shot medical landmark detection task. The x-axis refers to MRE and the y-axis refers to the percentage of MRE lying in the corresponding ranges. **Evidently, the choice of template affects the performance significantly.**

labeled data. The aggregated results work as more reliable pseudo labels for unlabeled data in rest part of the method. Another group of researchers, aiming to achieve a high performance at a low labeling cost, propose a strategy to select instances for annotation incrementally [15,31,35,45]. The basic idea is to first train a model with few labeled data, and then use the model to select instances from unlabeled data iteratively, which are annotated by specialists for the next round of training. Meanwhile, some researchers attempt to drain all potential of limited labeled data. With the power of self-training and self-supervised learning [1,4,24,39,43,47,51], it is possible to develop a robust, few-shot model even with several labeled samples. For example, Yao *et al.* [42] introduce a self-supervised proxy task that matches multi-layers features from images with different augmentations in the training stage, and use a *single image* as the template, whose patches centered at landmarks are matched with target images to make predictions.

However, during our research following the work of [42], we observe an interesting phenomenon (see Figure 1). The template choice highly impacts the final performance.

The mean radial error (MRE) of our trained model varies from  $2.9\text{mm}$  under the “best” template to  $4.5\text{mm}$  under the “worst” template. It is evident that there is a large gap lying between the best and the worst choices. Thus, a **selection question** naturally stands out: *Regarding the “gap” over samples, how to find and annotate the most “valuable” images in order to achieve the best performance with a deep model trained under such a limited supervision?* To the best of our knowledge, there is no ready answer to the above question. In this paper, we attempt to fill this blank.

To answer this question, we have to address three difficulties. (1) **No supervised signal**: For a landmark detection task, there are no labels to guide our image selection — We need to find substitutes for landmarks; (2) **No proper metric**: Mean radial error (MRE) is often employed as a performance metric for landmark detection, but we cannot compute MRE when no landmarks are available — we need to find proxy for MRE; (3) **Missing effective feature extraction**: Can we train a deep model that effectively extracts the features for template selection?

In this paper, we propose a framework named Sample Choosing Policy (SCP) to find the most annotation-worthy images as templates. First, to handle the situation of no landmark label, we choose handcrafted key points as substitutes for landmarks of interest. Second, to replace the MRE, we proposed to use a similarity score between a template and the rest based on the features of such potential key points. Third, considering landmark detection is a pixel-wise task, we apply pixel-wise self-supervised learning with all non-labeled data to build a basic feature extractor, which extracts features from each image. Finally, we can find out the subset of images with the highest similarities as the candidate templates to be labeled, from which a model is learned for few-shot landmark detection. With the help of SCP, we improve the MRE performance of one-shot medical landmark detection from  $3.595\text{mm}$  (with a random template) to  $3.083\text{mm}$  (with our selected template) in Cephalometric Xray dataset and  $4.114\text{mm}$  (with a random template) to  $2.635\text{mm}$  (with our selected template) in Hand Xray dataset; refer to Section 4.

## 2. Related Work

### 2.1. Active Learning

Active Learning queries a user interactively to label picked examples, while maximizing the performance gain of neural networks [10, 29, 33, 35, 45, 50]. In active learning, the model is first initialized by a subset of annotated examples, and then a certain strategy is suggested to select samples to label from the unlabeled data pool. The decision of selecting a sample depends on the information gained from the model, which is measured by entropy, divergence, gradients, etc. For example, the decision-making

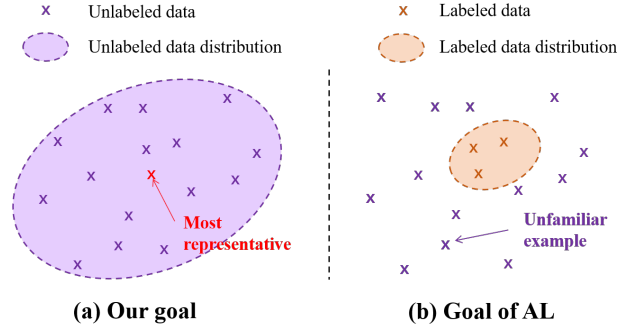


Figure 2. **Difference from active learning.** Deep models can remember and cluster the images or patterns they viewed. Instead of active learning (AL) tending to find the unfamiliar examples, our goal is to find the ones nearest to the center of latent space which we think more representative and important when only several images can be labeled.

strategy in [31, 35] is formed by a discriminator, assisted by a variational auto-encoder (VAE) that encodes the distribution of labeled data in a latent space. Kim et al [15] improve the strategy from a task perspective by updating it from a task-agnostic to task-aware style. More precisely, every selection in active learning aims to find hard examples in unlabeled dataset. However, the aforementioned selection question should be answered by a specific partition of samples, whose distribution is ideally as close as possible to the entire dataset, not just unfamiliar examples (Figure 2). In addition, some active learning methods like [29] also aim to collect core samples but require a labeled subset as a prerequisite. Given that, active learning is not the key to solving the selection problem.

### 2.2. Self-supervised Learning

Self-supervised learning (SSL) leverages information in data itself as the supervision, providing a solution for training from unlabeled data. Among all varieties of SSL, contrastive learning (CL) is one of the most powerful paradigms, leading to state-of-the-art performances in many vision tasks. Contrastive learning aims to pull the representations of similar samples closer, and dissimilar ones far apart. Most existing methods, like MoCo [7, 12], SimCLR [5, 6], BYOL [9] and BarlowTwins [44], are designed and optimized in instance-level comparisons, benefit the trained model with more discriminative and generalized representations. But it leads to sub-optimal representations for downstream tasks requiring pixel-level prediction, e.g., segmentation, object detection, and landmark detection. Recently there are some pixel-level methods attempting to learn dense feature representations. Multi-scale pixel-wise contrastive proxy task based on InfoNCE loss [11, 23] is introduced in [42] and achieves great performance in medical

landmark detection, which is used for our feature extractor.

### 2.3. Few-shot Learning

Neural networks are often hampered when there are only a few training samples to provide supervisory signals. Few-shot learning (FSL) [16, 21, 27, 52] is proposed to tackle this data scarcity problem, utilizing external prior knowledge to discover patterns in data. FSL has been employed successfully in tasks like segmentation [22, 46], detection [8, 32], etc. A typical example of FSL is the pretrain-finetune paradigm. For example, Bjorn *et al.* [2] leverage image reconstruction to pretrain a model, and finetune it with a few samples for the target facial landmark task. For medical landmark detection, some researchers treat the few labeled samples as templates. They start with a pre-trained model from other unlabeled data and make the predictions by matching target instances and templates. Yao *et al.* [42] introduce a pixel-wise multi-layer proxy task for self-supervision, which provides more generalized features utilized in the template-matching process. Lei *et al.* [18] propose another proxy task of predicting the relative offsets of two patches from an identical image. Not only that, the relative offsets also work as a metric in template matching.

## 3. Method

In this section, we introduce the proposed framework named Sample Choosing Policy (SCP) in detail.

### 3.1. Sample choosing policy

**Landmark detection using template matching.** Given a template  $T$ , landmark detection for an image  $X$  is first implemented via classical template matching.

Denote the set of landmark by  $P = \{p_1, p_2, \dots, p_L\}$ . Suppose that  $p_l^T \in P^T$  is the  $l^{th}$  landmark point in the template  $T$ , its corresponding landmark  $p_l^X$  in the image  $X$  is found by the following *searching-and-maximizing* problem:

$$p_l^X = \arg \max_p s[ F_\theta \circ T(p_l^T), F_\theta \circ X(p) ]; p_l^T \in P^T, \quad (1)$$

where  $p$  is coordinates of a pixel,  $s$  is a similarity function,  $F_\theta$  is a feature extractor, and  $F_\theta \circ X(p)$  computes the feature map for the image  $X$  and then extracts the feature vector at pixel  $p$ . The maximum value of the similarity function  $s$  achieved by  $p_l^X$  is denoted by  $r_l[T \rightarrow X]$ :

$$r_l[T \rightarrow X] = s[ F_\theta \circ T(p_l^T), F_\theta \circ X(p_l^X) ]. \quad (2)$$

The above landmark detection process considers only one template. When there are multiple templates  $\{T_1, T_2, \dots, T_M\}$  indexed by  $m = 1 : M$ , template matching is implemented by

$$(m_l, p_l^X) = \arg \max_{(m,p)} s[ F_\theta \circ T_m(p_l^T), F_\theta \circ X(p) ], \quad (3)$$

which finds the best template  $T_{m_l}$  for each landmark  $l$  as well as the matched landmark location  $p_l^X$ . By the same token, the maximum value of the similarity function  $s$  achieved by  $(m_l, p_l^X)$  is denoted by  $\hat{r}_l[\{T_m\} \rightarrow X]$ :

$$\hat{r}_l[\{T_m\} \rightarrow X] = s[ F_\theta \circ T_{m_l}(p_l^T), F_\theta \circ X(p_l^X) ]. \quad (4)$$

For the choice of *similarity function*, we utilize the commonly used cosine similarity function:

$$s[v^T, v^X] = \text{CosSim}(v^T, v^X) = \frac{\langle v^T \cdot v^X \rangle}{\|v^T\|_2 \cdot \|v^X\|_2}, \quad (5)$$

where  $v$  is a feature vector.

**Template selection.** To single out the best  $M$  templates among a set of images  $\Omega = \{X_1, X_2, \dots, X_N\}$ , we aim to seek the set of templates  $\{T_m\}$  that contains “the most similar” landmark information for all landmarks and with respect to all images:

$$\{\hat{T}_m\} = \arg \max_{\{T_m\} \subset \Omega} \frac{1}{N} \sum_n \frac{1}{L} \sum_l \hat{r}_l[\{T_m\} \rightarrow X_n]. \quad (6)$$

The above optimization is *combinatorial in nature* as there are  $\binom{N}{M}$  possible combinations, which are nearly impossible to exhaust in practice except for very small  $M$ . Therefore, we randomly sample a large number (say 10,000) of combinations and pick the maximizing combination as an approximate solution.

**Substituting landmarks with key points.** To implement template selection per Eq. (6), the knowledge of landmarks is assumed. However, even such knowledge is nonexistent before template selection. Therefore, we proposed to utilize potential key points to substitute landmarks. In particular, we utilize the classical multi-scale detector, SIFT, to find key points, where landmarks are likely to co-locate.

For each image  $X \in \Omega$ , we apply SIFT to get its corresponding  $K$  key points  $Q^X = \{q_1^X, q_2^X, \dots, q_K^X\}$  with the highest responses. Further, the SIFT key points for different images are not in correspondence, directly applying Eq. (3) is not possible. To address such an issue, we perform the template matching in a *reverse order*, that is, for an image  $X$  with its key points  $Q^X$ , we perform the following for each key point  $q_k^X$ :

$$(m_k, q_k^T) = \arg \max_{(m,q)} s[ F_\theta \circ T_m(q), F_\theta \circ X(q_k^X) ], \quad (7)$$

and record the achieved maximum as

$$\hat{r}_k[X \rightarrow \{T_m\}] = s[ F_\theta \circ T_{m_k}(q_k^T), F_\theta \circ X(q_k^X) ]. \quad (8)$$

Finally, we define the average similarity  $R$  as the representative score of  $\{T_m\}$ :

$$R[\{T_m\}] = \frac{1}{N} \sum_n \frac{1}{K} \sum_k \hat{r}_k[X_n \rightarrow \{T_m\}]. \quad (9)$$

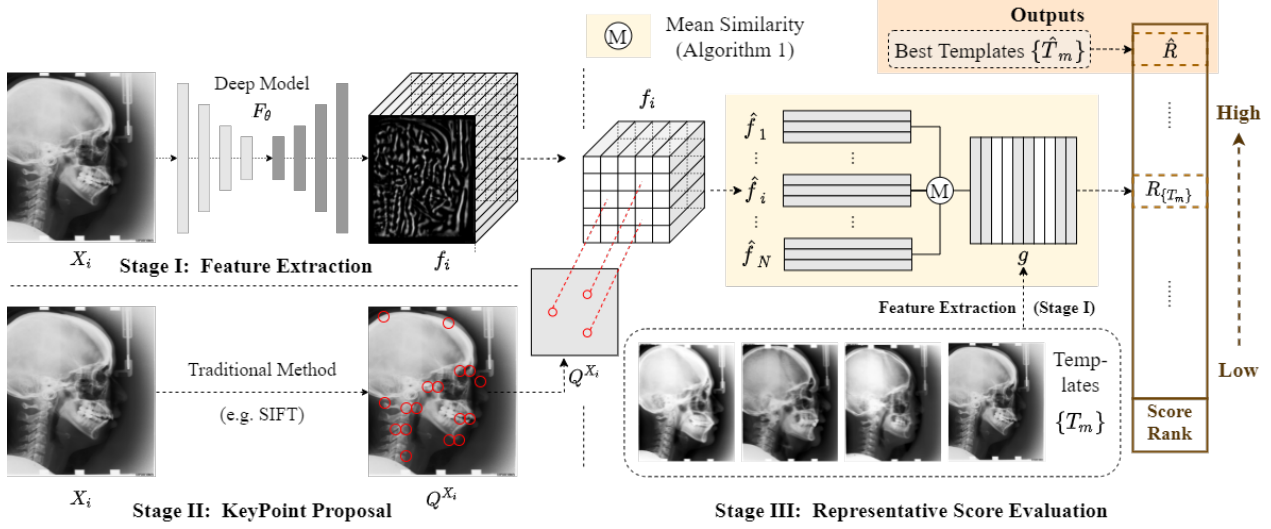


Figure 3. **Overview of Sample Choosing Policy (SCP):** SCP consists of three stages. Stage 1: Extract features  $f_i$  from image  $X_i$  via a self-trained deep model. Stage 2: Detect key points  $Q^{X_i}$  by traditional method like SIFT from image  $X_i$ . Stage 3: Evaluate similarities between filtered features  $\hat{f}_i$  and features  $g$  of templates  $\{T_m\}$ , and record the average similarity  $R[\{T_m\}]$ . The combination of templates with highest similarity  $\hat{R}$  are selected as  $\{\hat{T}_m\}$ .

and Eq. (6) is accordingly adapted:

$$\{\hat{T}_m\} = \arg \max_{\{T_m\} \subset \Omega} R[\{T_m\}] \quad (10)$$

**Construction of feature extractor using contrastive learning.** To answer the question that what kind of deep model can support us to make selection, we start our analysis of Eq. (1), which maximizes the similarities between the same landmarks from different images.

Without landmark labeling, we resort to contrastive learning, which is proven to be a reliable tool to learn a basic model without using any label information. Here, instead of instance-level self-supervised learning for visual recognition [9, 44], we adapt it for achieving the goal of (1). We do so by narrowing the distance between different views of an identical patch and extend the distance between different patches. For example, InfoNCE loss [23] is widely used and applied in our training for feature extractor,

$$\begin{aligned} \mathcal{L}_{\text{InfoNCE}} &= -\mathbb{E} \left[ \log \frac{\exp(\alpha)}{\exp(\alpha) + \sum \exp(\alpha')} \right]; \\ \alpha &= s[F_\theta \circ X(p), F_\theta \circ X_{\text{aug}}(p)]; \\ \alpha' &= s[F_\theta \circ X(p), F_\theta \circ X(q)], \end{aligned} \quad (11)$$

where  $X_{\text{aug}}$  is a different version of  $X$  by augmentation, and  $p$  and  $q$  are two different key points.

We follow [41, 42] to construct a deep model trained via multi-layer pixel-wise contrastive loss function as our feature extractor. According to [42], we use VGG [30] as the backbone, followed with 5 blocks to reduce the dimension.

This model is trained with a pixel-wise matching proxy task for over 500 epochs.

**The overall template selection pipeline.** Based on the above discussion, our pipeline is summarized as in Figure 3, assuming the availability of the feature extractor  $F_\theta$ , which is learned using the self-supervised pixel-wise matching task. First, we extract features from all images  $\Omega$  and candidate templates  $\{T_m\}$  for the following operations. Second, we extract key points  $Q^X$  with the help of traditional key point detector SIFT. Third, we pair each image  $X_n$  and the template group  $\{T_m\}$  to obtain the similarity  $r_n$  between  $X_n$  and  $\{T_m\}$  (Eq. (8)). The mean similarity  $R_{\{T_m\}}$  of all pairs of  $X_n$  and  $\{T_m\}$  indicate the “representativeness” of  $\{T_m\}$  to the whole dataset, and the best group of templates  $\{\hat{T}_m\}$  achieving the maximum of mean similarity are our final selection (Eq. (9)).

### 3.2. Refined landmark detection

While landmark detection is implemented as template matching in Section 3.1, its detection performance is still limited as its feature detector is geared for all pixels not specifically for the landmarks. We further follow [42] to improve the detection of landmarks via semi-supervised learning. Another landmark detection deep model with a heatmap predictor and two offset predictors (offsets in x- and y-axis) is built for distilling with pseudo landmark labels predicted by the previous template matching model. The performance of the distilled model is geared toward landmarks of interest and is better than the previous model.



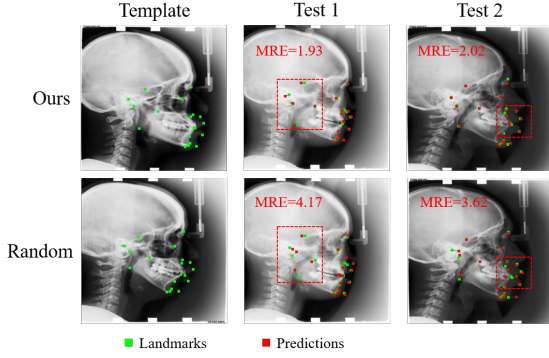


Figure 4. **Visual Comparison of templates from our policy and random selection.** Column “Template/Test 1/Test 2” refers to the templates and two test images. The row “Ours” and “Random” refers to the template selected by our method and random selection, respectively. As shown in red dashed boxes, our template outperforms the random selected template in visualization.

## 4. Experiments

### 4.1. Datasets

**Cephalometric Xray:** It is a widely-used public dataset for cephalometric landmark detection, containing 400 radiographs, and is provided in IEEE ISBI 2015 Challenge [14, 37]. There are 19 landmarks of anatomical significance labeled by 2 expert doctors in each radiograph. The averaged version of annotations by two doctors is set as the ground truth. The image size is  $1935 \times 2400$  and the pixel spacing is 0.1mm. The dataset is split into 150 and 250 for training and testing respectively, referring to the official division.

**Hand Xray:** It is also a public dataset including 909 X-ray images of hands. The setting of this dataset follows [25]. The first 609 images are used for training and the rest for testing. The image size varies among a small range, so all images are resized to  $384 \times 384$ .

**WFLW:** This dataset is from [38] containing 10,000 faces with 7500 and 2500 in training and test sets, respectively. All images are collected from the WIDER FACE dataset [40] and manually labeled with 98 landmarks. The dataset contains different test subsets where the image appearances vary due to variations in pose, expression, and/or illumination or the presence of blur, occlusion, and/or make-up.

### 4.2. Settings

**Metrics:** Following the official challenge [14, 37], mean radial error (MRE) and successful detection rate (SDR) in four radii (2mm, 2.5mm, 3mm, and 4mm) are applied, based on the Euclidean distance between prediction and ground truth. In addition, similarity via Eq. (9) is demonstrated for comparison. For the WFLW dataset, the performance of facial landmark detection is reported by normal-

ized mean error (NME), failure rate at 10% NME (FR@10), and area-under-the-curve (AUC) of the Cumulative Error Distribution (CED) curve.

**Implementation details:** All of our models are implemented in PyTorch, accelerated by an NVIDIA RTX GPU. Following [41, 42], the feature extractor is optimized by Adam optimizer for 3500 epochs for self-supervised training. The learning rate is initialized with 0.001, and decayed by half every 500 epochs. It takes 6 hours to converge with batch size set to 8.

### 4.3. The performance of SCP

**Few-shot medical landmark detection:** Firstly, experiments are conducted on different numbers of templates for Cephalometric dataset. For Table 1,  $M$  denotes the number of templates used in experiment. The columns “ours” refer to the results achieved by the proposed method, while “random” means to the average results of multiple rounds of training (we use 1,000 trials), with standard variations. The column “best” and “worst” refer to the best or worst results in multiple tries, respectively.

As shown in Table 1, our approach cannot find the best templates, instead, it finds a fairly good choice well above the average performance. When there is only one template, the best template achieves 2.863mm in MRE. Although the template we choose achieves an MRE of only 3.083mm, it is much better than the average MRE 3.595mm. Another fact revealed by Table 1 is the “diminishing return” phenomenon: As the number of templates increases, the difference of maximum similarities and performances between choices tends to be smaller. In our experiment, we find that when  $M$  achieves 75 (half of the dataset) or more, our template selection policy has little effect because of sufficient modeling of the appearance variance in medical images.

The predictions by one template from our method and random selection are visualized in Figure 4. Our predictions around ears and nose locate more closer to the ground-truth landmarks than those by random template, which has consistent performance in MRE, quantitatively.

Besides, another group of experiments are conducted on Hand Xray dataset. The proposed SCP is applied on Hand Xray dataset to obtain the suggested  $M$  templates. Following settings in [42], evaluation model is built. Results are listed in Table 1, showing reliable improvements (e.g., MRE reduced by 35.5% (4.114mm to 2.653mm)). As reported above, MRE results of 5/10/15 templates perform a bit worse than that of 1 template (e.g., the MRE is 2.891mm for 5 templates, but 2.653mm for 1 template). The dataset size of Hand Xray (609 images) is much bigger than that of Cephalometric (150 images). We speculate that the diversity of Hand Xray dataset could not be well “represented” by such small group of templates. The smaller the number of templates is, the more randomness it underlies. So it

Table 1. Comparison of different numbers of templates on the ISBI 2015 Challenge [37] and Hand Xray [42] testsets. The landmarks are detected based template matching.

Dataset	M	Similarity ( $\uparrow$ )		MRE ( $\downarrow$ ) (mm)			
		ours	random	ours	random	best	worst
Cephalometric	1	0.554	0.519 $\pm$ 0.017	3.083	3.595 $\pm$ 0.381	2.863	4.952
	2	0.638	0.604 $\pm$ 0.012	2.840	3.177 $\pm$ 0.227	2.661	4.090
	3	0.676	0.646 $\pm$ 0.009	2.742	3.019 $\pm$ 0.187	2.573	3.875
	4	0.699	0.672 $\pm$ 0.008	2.677	2.912 $\pm$ 0.164	2.512	3.674
	5	0.709	0.690 $\pm$ 0.007	2.571	2.850 $\pm$ 0.141	2.486	3.458
Hand Xray	1	0.363	0.288 $\pm$ 0.052	2.635	4.114 $\pm$ 3.566	2.468	43.81
	5	0.532	0.482 $\pm$ 0.021	2.891	3.188 $\pm$ 0.493	2.377	6.268
	10	0.601	0.547 $\pm$ 0.013	2.769	3.070 $\pm$ 0.346	2.289	4.808
	15	0.615	0.581 $\pm$ 0.011	2.664	2.989 $\pm$ 0.311	2.259	4.245
	50	0.690	0.673 $\pm$ 0.006	2.524	2.740 $\pm$ 0.219	2.240	3.775

Table 2. Comparison of the state-of-the-art landmark detection approaches and our refined approach on the ISBI 2015 Challenge [37] testset.

Model	Labeled images	MRE ( $\downarrow$ ) (mm)	SDR ( $\uparrow$ ) (%)			
			2mm	2.5mm	3mm	4mm
Ibragimov <i>et al.</i> [13]*	150	-	68.13	74.63	79.77	86.87
Lindner <i>et al.</i> [20]*	150	1.77	70.65	76.93	82.17	89.85
Urschler <i>et al.</i> [36]*	150	-	70.21	76.95	82.08	89.01
Payer <i>et al.</i> [26]*	150	-	<b>73.33</b>	<b>78.76</b>	<b>83.24</b>	<b>89.75</b>
Payer <i>et al.</i> [26] <sup>†</sup>	25	2.54	<b>66.12</b>	<b>73.27</b>	<b>79.82</b>	86.82
Yao <i>et al.</i> [42] <sup>†</sup>	25	2.17	55.83	66.65	76.88	89.13
SCP (Ours w/ refinement)	25	<b>2.06</b>	58.00	68.52	79.13	<b>90.40</b>
Payer <i>et al.</i> [26] <sup>†</sup>	10	6.52	49.49	57.91	65.87	75.07
Yao <i>et al.</i> [42] <sup>†</sup>	10	2.30	53.11	63.58	74.34	87.17
SCP (Ours w/ refinement)	10	<b>2.24</b>	<b>55.41</b>	<b>64.67</b>	<b>75.26</b>	<b>87.38</b>
Payer <i>et al.</i> [26] <sup>†</sup>	5	12.34	27.35	32.94	38.48	45.28
Yao <i>et al.</i> [42] <sup>†</sup>	5	2.44	49.92	60.83	71.76	85.13
SCP (Ours w/ refinement)	5	<b>2.33</b>	<b>52.25</b>	<b>63.05</b>	<b>73.95</b>	<b>86.27</b>
Yao <i>et al.</i> [42] <sup>†</sup>	1	2.90	37.16	48.04	60.02	77.72
SCP (Ours w/ refinement)	1	<b>2.74</b>	<b>43.79</b>	<b>53.05</b>	<b>64.12</b>	<b>79.05</b>

\* copied from their original papers.

<sup>†</sup> re-implemented with limited labeled images.

makes sense the results tend to be stable when the number of templates increases.

**Refined landmark detection results:** To compare with supervised landmark detection methods, we follow [42] and apply semi-supervised refinement based on the self-supervised results to achieve better performance (refer to Section 3.2). Our methods and [42] are trained with labeled and unlabeled data, while [13, 20, 26, 36] are trained with labeled images. As in Table 2 our method outperforms others when labels are limited (e.g., the MRE is 2.33mm for SCP with 5 labeled images, 2.44mm for [42], and 12.34mm for [26]), and achieves detection results comparable to the supervised methods that are trained based on even more labeled images (e.g., the MRE for [26] with 25 labeled images

is 2.54mm).

**Facial landmark detection:** We attempt to apply our templates selection method for [2] to improve the performance on facial landmark detection. The pre-trained model from the first stage of [2] is integrated into the proposed framework, which suggest the  $M$  instances with the highest mean similarities ( $M=50$  in our experiment). In order to eliminate unstable factors, models with the first  $M$  instances or instances from our policy are both trained for 5 times. The evaluation is conducted on full and make-up test split of WFLW dataset. The results are listed in Table 3 where our method significantly improves NME from 9.131 to 8.998 in full split and from 8.724 to 8.578 in make-up split.

Table 3. Improvements on WFLW dataset based on [38] with 50 labeled images.

WFLW dataset			
Full	NME( $\downarrow$ )	FR@10( $\downarrow$ )	AUC( $\uparrow$ )
Original <sup>†</sup>	9.131 $\pm$ 0.316	27.506 $\pm$ 1.785	0.278 $\pm$ 0.021
Ours	<b>8.998<math>\pm</math>0.129</b>	<b>27.475<math>\pm</math>0.840</b>	<b>0.285<math>\pm</math>0.002</b>
Make-up	NME( $\downarrow$ )	FR@10( $\downarrow$ )	AUC( $\uparrow$ )
Original <sup>†</sup>	8.724 $\pm$ 0.054	25.79 $\pm$ 1.689	0.283 $\pm$ 0.003
Ours	<b>8.578<math>\pm</math>0.188</b>	<b>25.79<math>\pm</math>0.989</b>	<b>0.284<math>\pm</math>0.003</b>

<sup>†</sup> re-implemented from [2]

## 5. Ablation Study

### 5.1. Key Point Proposal

*Q: How good is the use of SIFT key points as substitutes for landmarks?* Figure 5 demonstrate the relationship between landmarks and potential key points from handcraft methods in feature level (Eq. (9)). The similarities calculated by potential key points are positively correlated with those by landmarks to a large extent. Thus, we pick the potential key points as the substitutes. It should be noted that excessive key points will not improve performance but increase the computational cost (e.g. 2.571mm for 100 key points and 2.599mm for 150 key points in Table 5), but sparse key points will lead to performance degradation (e.g. MRE increase from 2.571mm to 2.670mm as the number of key points decrease from 100 to 50.)

*Q: Do other kinds of handcrafted key points also work in our framework?* In the proposed method, we use SIFT points as the potential key points to avoid estimating features of all pixels. To compare with other key point detection methods, we adopt SIFT, SURF, ORB, and random selection as our key point selectors. For SIFT, SURF, and ORB, we detect the key points and then filter out some very close points. For random selection, we randomly select 100 points as the key points. The results are listed on Table 4, where we find the performance of SURF is close to SIFT, and both greatly exceed the random filtering (e.g., SIFT/SURF are both 2.571mm for 5 templates, and greater than 2.835mm of random points). The locations of ORB key points concentrate in several area and performs less robust than SIFT and SURF.

### 5.2. Selection Policy

*Q: How good is the proposed similarity correlated with MRE?* The metric of MRE is widely used for evaluating the landmark detection performance. Therefore, it is of interest to know how the proposed similarity is correlated with MRE. Figure 6 presents such a plot, which nicely shows a strong (negative) correlation coefficient of  $-0.675$ . This enables us to utilize the proposed similarity function as a




Figure 5. **Similarities of potential key points vs. landmarks.** The correlation coefficient (CC) of potential key points and landmarks is 0.462, thus we think it is feasible to replace landmarks with potential key points when estimating similarities.




Figure 6. **Mean Similarity vs. Mean Radial Error (MRE).** Similarities are estimated by calculating the average similarities of all landmarks. The correlation coefficient (CC) of the two indicators is  $-0.675$ , indicating a strong, negative correlation.

proxy of MRE, which is more convenient for us to estimate when there is no label.

*Q: What about other selection methods?* We also try some other methods based on ideas from active learning. VAAL [31] is a typical active learning framework which we re-implement and start with random 5 instances. Next, we obtain  $N(N = 5)$  instances suggested by VAAL to be templates. Uncertainty and Entropy are also common tools for active learning [28]. We re-implement the framework by connecting one encoder with two classifiers, and train by classifying all instances and enlarging the difference of outputs from two classifiers. Finally, we estimate the entropy

Table 4. Comparison for different SSL methods and different traditional key point detectors on Cephalometric dataset.

M	MRE ( $\downarrow$ ) (mm)						
	SSL methods			Key Point Detectors			
	Ours	BYOL-m [9]	BYOL [9]	SIFT	SURF	ORB	Random
1	<b>3.083</b>	3.366	3.362	<b>3.080</b>	3.208	3.668	3.181
2	<b>2.846</b>	3.086	3.085	<b>2.846</b>	2.901	3.296	3.219
3	<b>2.743</b>	3.008	2.877	<b>2.743</b>	2.787	2.896	2.860
4	<b>2.673</b>	2.894	2.761	<b>2.673</b>	2.621	2.621	2.887
5	<b>2.571</b>	2.762	2.739	<b>2.571</b>	2.571	2.677	2.835

Table 5. Ablation study of different number of SIFT key points and searches for selecting best templates on Cephalometric dataset.

Ablation	Num	Similarity ( $\uparrow$ )	MRE ( $\downarrow$ ) (mm)
Keypoints	50	0.7128	2.670
	75	0.7099	2.623
	100	<b>0.7117</b>	<b>2.571</b>
	150	0.7105	2.599
Searches	1000	0.7080	2.711
	10000	0.7105	2.703
	50000	0.7108	2.577
	100000	<b>0.7117</b>	<b>2.571</b>

Table 6. Comparison with different sampling methods on Cephalometric dataset.

Method	Similarity( $\uparrow$ )	MRE( $\downarrow$ ) (mm)
Ours	<b>0.709</b>	<b>2.571</b>
VAAL [31]	0.691 $\pm$ 0.004	2.851 $\pm$ 0.077
Loss(MSE)	0.695 $\pm$ 0.008	2.766 $\pm$ 0.093
Uncertainty	0.694 $\pm$ 0.005	2.876 $\pm$ 0.074
Entropy	0.692 $\pm$ 0.006	2.847 $\pm$ 0.094

(of one-hot vectors), uncertainty (difference of two classifiers) and loss as metrics to suggest templates to label. As shown in Table 6, VAAL needs labeled data to initialize, and performs badly in only one iteration. MSE works while uncertainty and entropy are not, where the probable reason is that patterns in medical images appear simple for both classifiers which give similar predictions and cause very low uncertainty for all images and low entropy to distinguish instances clearly.

*Q: How does the number of combinations affect the final performance?* As the number of templates increase, it becomes much harder to iterate all combinations. Therefore, we randomly choose the templates in limited iterations to find the best combination. We conduct an experiment about the number of random selections which is shown in Table 5. As the number of random selections increases, we have more chances to find the combination with a larger similarity, which yields better landmark detection perfor-

mance.

### 5.3. Self-supervised Method

*Q: Can self-supervised methods affect our performance?* In the proposed method, we leverage the proxy task in [42] to pre-train our feature extractor. In addition, we implement other self-supervised methods, including BYOL [9] in pixel level, and BYOL with multi-layer training (BYOL-m). From Table 4 we discover the strong correlation between the performance of self-supervised models and the quality of selection results and our model with multi-layer InfoNCE outperforms others in feature extraction (e.g. 3.083mm for ours and 3.366/3.263mm for others) due to more fitting the landmark detection task.

## 6. Conclusion and Future work

We propose Sample Choice Policy (SCP) for few-shot medical landmark detection task, a novel framework for screening out the representative instances to reduce labor on annotation and achieve high performance simultaneously. SCP learns to map the consistent anatomical information into feature spaces by solving a self-supervised proxy task and extracting representative patches from all images in the first stage. Simultaneously, SCP leverages traditional key point detector to pick out valuable patches with a large local variation or steep edges in images as the representatives. Finally, SCP estimates the relevance between images by averaging the similarities between their representative patches, and selects the images with high average relevance with all other images. Our extensive experiments show that SCP outperforms the conventional policies of template selection and, after refinement, achieves state-of-the-art performances for few-shot medical landmark detection. In the future, we plan to further improve the efficiency and accuracy of finding out the best templates by designing deep models that extract more representative features and to explore the idea of template selection for other applications such as image segmentation.



## References

- [1] Riddhish Bhalodia, Ladislav Kavan, and Ross T Whitaker. Self-supervised discovery of anatomical shape landmarks. In *MICCAI*, pages 627–638. Springer, 2020. [1](#)
- [2] Björn Browatzki and Christian Wallraven. 3fabrec: Fast few-shot face alignment by reconstruction. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 6109–6119. IEEE, 2020. [3](#), [6](#), [7](#)
- [3] Runnan Chen, Yuexin Ma, Lingjie Liu, Nenglun Chen, Zhiming Cui, Guodong Wei, and Wenping Wang. Semi-supervised anatomical landmark detection via shape-regulated self-training. *arXiv preprint arXiv:2105.13593*, 2021. [1](#)
- [4] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607. PMLR, 2020. [1](#)
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 1597–1607. PMLR, 2020. [2](#)
- [6] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E. Hinton. Big self-supervised models are strong semi-supervised learners. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [2](#)
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. [2](#)
- [8] Zhibo Fan, Yuchen Ma, Zeming Li, and Jian Sun. Generalized few-shot object detection without forgetting. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 4527–4536. Computer Vision Foundation / IEEE, 2021. [3](#)
- [9] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent - A new approach to self-supervised learning. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. [2](#), [4](#), [8](#)
- [10] Denis A. Gudovskiy, Alec Hodgkinson, Takuya Yamaguchi, and Sotaro Tsukizawa. Deep active learning for biased datasets via fisher kernel self-supervision. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 9038–9046. Computer Vision Foundation / IEEE, 2020. [2](#)
- [11] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In Yee Whye Teh and D. Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010*, volume 9 of *JMLR Proceedings*, pages 297–304. JMLR.org, 2010. [2](#)
- [12] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 9726–9735. Computer Vision Foundation / IEEE, 2020. [2](#)
- [13] Bulat Ibragimov, Boštjan Likar, F Pernus, and Tomaž Vrtovec. Computerized cephalometry by game theory with shape-and appearance-based landmark refinement. 2015. [6](#)
- [14] Kaggle. Cephalometric X-Ray Landmarks Detection Challenge, 2015. <https://www.kaggle.com/jiahongqian/cephalometric-landmarks/discussion/133268>. [5](#)
- [15] Kwanyoung Kim, Dongwon Park, Kwang In Kim, and Se Young Chun. Task-aware variational adversarial active learning. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 8166–8175. Computer Vision Foundation / IEEE, 2021. [1](#), [2](#)
- [16] Nameyong Kwon, Hwidong Na, Gabriel Huang, and Simon Lacoste-Julien. Repurposing pretrained models for robust out-of-domain few-shot learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [3](#)
- [17] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. [1](#)
- [18] Wenhui Lei, Wei Xu, Ran Gu, Hao Fu, Shaoting Zhang, and Guotai Wang. Contrastive learning of relative position regression for one-shot object localization in 3d medical images, 2021. [3](#)
- [19] Han Li, Long Chen, Hu Han, and S. Kevin Zhou. Conditional training with bounding map for universal lesion detection. *CoRR*, abs/2103.12277, 2021. [1](#)
- [20] Claudia Lindner and Tim F Cootes. Fully automatic cephalometric evaluation using random forest regression-voting. *Scientific Reports*, 6:33581, 2016. [6](#)
- [21] Lu Liu, William L. Hamilton, Guodong Long, Jing Jiang, and Hugo Larochelle. A universal representation transformer layer for few-shot image classification. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [3](#)
- [22] Khoi Nguyen and Sinisa Todorovic. FAPIS: A few-shot anchor-free part-based instance segmenter. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*

- 2021, virtual, June 19-25, 2021, pages 11099–11108. Computer Vision Foundation / IEEE, 2021. [3](#)
- [23] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [2](#), [4](#)
- [24] Cheng Ouyang, Carlo Biffi, Chen Chen, Turkay Kart, Huaqi Qiu, and Daniel Rueckert. Self-supervision with superpixels: Training few-shot medical image segmentation without annotation. In *ECCV*, pages 762–780. Springer, 2020. [1](#)
- [25] Christian Payer, Darko Štern, Horst Bischof, and Martin Urschler. Integrating spatial configuration into heatmap regression based cnns for landmark localization. *Medical Image Anal.*, 54:207–219, 2019. [5](#)
- [26] Christian Payer, Darko Štern, Horst Bischof, and Martin Urschler. Integrating spatial configuration into heatmap regression based cnns for landmark localization. *Medical image analysis*, 54:207–219, 2019. [6](#)
- [27] Mengye Ren, Michael Louis Iuzzolino, Michael Curtis Mozer, and Richard S. Zemel. Wandering within a world: Online contextualized few-shot learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. [3](#)
- [28] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Xiaojiang Chen, and Xin Wang. A survey of deep active learning. *CoRR*, abs/2009.00236, 2020. [7](#)
- [29] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. [2](#)
- [30] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. [4](#)
- [31] Samarth Sinha, Sayna Ebrahimi, and Trevor Darrell. Variational adversarial active learning. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5971–5980. IEEE, 2019. [1](#), [2](#), [7](#), [8](#)
- [32] Bo Sun, Banghuai Li, Shengcai Cai, Ye Yuan, and Chi Zhang. FSCE: few-shot object detection via contrastive proposal encoding. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 7352–7362. Computer Vision Foundation / IEEE, 2021. [3](#)
- [33] Iris Sundin, Peter Schulam, Eero Siivola, Aki Vehtari, Suchi Saria, and Samuel Kaski. Active learning for decision-making from imbalanced observational data. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6046–6055. PMLR, 2019. [2](#)
- [34] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*, 2017. [1](#)
- [35] Toan Tran, Thanh-Toan Do, Ian D. Reid, and Gustavo Carneiro. Bayesian generative active deep learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 6295–6304. PMLR, 2019. [1](#), [2](#)
- [36] Martin Urschler, Thomas Ebner, and Darko Štern. Integrating geometric configuration and appearance information into a unified framework for anatomical landmark localization. *Medical image analysis*, 43:23–36, 2018. [6](#)
- [37] Ching-Wei Wang, Cheng-Ta Huang, Jia-Hong Lee, Chung-Hsing Li, Sheng-Wei Chang, Ming-Jhih Siao, Tat-Ming Lai, Bulat Ibragimov, Tomaž Vrtovec, Olaf Ronneberger, et al. A benchmark for comparison of dental radiography analysis algorithms. *Medical Image Analysis*, 31:63–76, 2016. [5](#), [6](#)
- [38] Wayne Wu, Chen Qian, Shuo Yang, Quan Wang, Yici Cai, and Qiang Zhou. Look at boundary: A boundary-aware face alignment algorithm. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 2129–2138. IEEE Computer Society, 2018. [5](#), [7](#)
- [39] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10687–10698, 2020. [1](#)
- [40] Wei Yang, Shuang Li, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. Learning feature pyramids for human pose estimation. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 1290–1299. IEEE Computer Society, 2017. [5](#)
- [41] Qingsong Yao, Zecheng He, Hu Han, and S Kevin Zhou. Miss the point: Targeted adversarial attack on multiple landmark detection. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 692–702. Springer, 2020. [4](#), [5](#)
- [42] Qingsong Yao, Quan Quan, Li Xiao, and S. Kevin Zhou. One-shot medical landmark detection. *CoRR*, abs/2103.04527, 2021. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [43] Qingsong Yao, Li Xiao, Peihang Liu, and S Kevin Zhou. Label-free segmentation of covid-19 lesions in lung ct. *IEEE Transactions on Medical Imaging*, 2021. [1](#)
- [44] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 12310–12320. PMLR, 2021. [2](#), [4](#)
- [45] Beichen Zhang, Liang Li, Shijie Yang, Shuhui Wang, Zheng-Jun Zha, and Qingming Huang. State-relabeling adversarial active learning. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 8753–8762. Computer Vision Foundation / IEEE, 2020. [1](#), [2](#)

- [46] Na Zhao, Tat-Seng Chua, and Gim Hee Lee. Few-shot 3d point cloud semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 8873–8882. Computer Vision Foundation / IEEE, 2021. [3](#)
- [47] Hong-Yu Zhou, Shuang Yu, Cheng Bian, Yifan Hu, Kai Ma, and Yefeng Zheng. Comparing to learn: Surpassing imagenet pretraining on radiographs by comparing image representations. In *MICCAI*, pages 398–407. Springer, 2020. [1](#)
- [48] S. Kevin Zhou, Hayit Greenspan, Christos Davatzikos, James S. Duncan, Bram van Ginneken, Anant Madabhushi, Jerry L. Prince, Daniel Rueckert, and Ronald M. Summers. A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. *Proceedings of the IEEE*, 2021. [1](#)
- [49] S Kevin Zhou, Daniel Rueckert, and Gabor Fichtinger. *Handbook of Medical Image Computing and Computer Assisted Intervention*. Academic Press, 2019. [1](#)
- [50] Zongwei Zhou, Jae Y. Shin, Lei Zhang, Suryakanth R. Gurudu, Michael B. Gotway, and Jianming Liang. Fine-tuning convolutional neural networks for biomedical image analysis: Actively and incrementally. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4761–4772. IEEE Computer Society, 2017. [2](#)
- [51] 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. [1](#)
- [52] Kai Zhu, Yang Cao, Wei Zhai, Jie Cheng, and Zheng-Jun Zha. Self-promoted prototype refinement for few-shot class-incremental learning. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 6801–6810. Computer Vision Foundation / IEEE, 2021. [3](#)