



Free Lunch for Surgical Video Understanding by Distilling Self-supervisions

Xinpeng Ding¹, Ziwei Liu², and Xiaomeng Li^{1,3}(✉)

¹ The Hong Kong University of Science and Technology, Kowloon, Hong Kong
eexmli@ust.hk

² S-Lab, Nanyang Technological University, Singapore, Singapore

³ The Hong Kong University of Science and Technology Shenzhen Research Institute,
Kowloon, Hong Kong

Abstract. Self-supervised learning has witnessed great progress in vision and NLP; recently, it also attracted much attention to various medical imaging modalities such as X-ray, CT, and MRI. Existing methods mostly focus on building new pretext self-supervision tasks such as reconstruction, orientation, and masking identification according to the properties of medical images. However, the publicly available self-supervision models are not fully exploited. In this paper, we present a powerful yet efficient self-supervision framework for surgical video understanding. Our key insight is to distill knowledge from *publicly available models trained on large generic datasets* (For example, the released models trained on ImageNet by MoCo v2: <https://github.com/facebookresearch/moco>) to facilitate the self-supervised learning of surgical videos. To this end, we first introduce a semantic-preserving training scheme to obtain our teacher model, which not only contains semantics from the publicly available models, but also can produce accurate knowledge for surgical data. Besides training with only contrastive learning, we also introduce a distillation objective to transfer the rich learned information from the teacher model to self-supervised learning on surgical data. Extensive experiments on two surgical phase recognition benchmarks show that our framework can significantly improve the performance of existing self-supervised learning methods. Notably, our framework demonstrates a compelling advantage under a low-data regime. Our code is available at <https://github.com/xmed-lab/DistillingSelf>.

Keywords: Surgical videos · Self-supervised learning · Knowledge distillation

1 Introduction

Generally, training deep neural networks requires a large amount of labeled data to achieve great performance. However, obtaining the annotation for surgical videos is expensive as it requires professional knowledge from surgeons. To



Fig. 1. Left: Existing self-supervised learning approaches aim to design different pretext tasks on a larger dataset to improve the representation ability. **Right:** We proposed a novel distilled contrastive learning to transfer the powerful representation ability from the free available models to the surgical video self-training.

address this problem, self-supervised learning, *i.e.*, training the model on a large dataset (*e.g.*, ImageNet [20]) without annotations and fine-tuning the classifier on the target datasets, has achieved great success for image recognition and video understanding [2, 4, 10, 19]. This paper focuses on designing self-supervised learning methods for surgical video understanding with a downstream task - surgical phase recognition, which aims to predict what phase is occurring for each frame in a video [1, 6, 9, 13–15, 24, 25, 28, 29].

Self-supervised learning has been widely applied into various medical images, such as X-ray [30], fundus images [16, 17], CT [34] and MRI [30, 32]. For example, Zhuang *et al.* [34] developed a rubik’s cube playing pretext task to learn 3D representation for CT volumes. Rubik’s cube+ [33] introduced more pretext tasks, *i.e.*, cube ordering, cube orientation and masking identification to improve the self-supervised learning performance. Some researchers introduced reconstruction of corrupted images [3, 32] or triplet loss [26] for self-supervised learning on nuclei images. Recently, motivated by the great success of contrastive learning in computer vision [2, 4, 5, 10, 23], Zhou *et al.* [31] introduced the contrastive loss to 2D radiographs. Taleb *et al.* [22] further improved contrastive predictive coding [18] to a 3D manner for training with 3D medical images. To learn more detailed context, PCRL [30] combined the construction loss with the contrastive one. However, all of these above approaches aim to design different pretext tasks to perform self-supervised learning on medical datasets, as shown in Fig. 1(a).

In this paper, we make a crucial observation that *using the same backbone and self-supervised learning method, the model trained with ImageNet data can yield a comparable performance for surgical phase recognition with that trained with surgical video data (82.3% vs 85.7% Acc.)*; see details in Table 1. This surprising result indicates that the model self-supervised trained with ImageNet data can learn useful semantics that benefit surgical video understanding. There are many publicly available self-supervised models, motivating us to leverage the free knowledge to facilitate the self-supervised learning of the surgical video.

To this end, we propose a novel distilled self-supervised learning method for surgical videos, which distillates the free knowledge from *publicly available self-supervised models (e.g., MoCo v2 trained on the ImageNet)* to improve the self-supervised learning on surgical videos, as shown in Fig. 1(b). We first introduce

a semantic-preserving training to train a teacher model, which retains representation ability from ImageNet while producing accurate semantics knowledge for surgical data. Besides training the student model with the only contrastive objective, we boost the self-supervised learning on the surgical data with a distillation objective. The proposed distillation objective forces the similarity matrix of the teacher and the student models to be consistent, which makes the student model to learn extra semantics.

We summarize our key contributions as follows:

- To best of our knowledge, we are the first to investigate the use of self-supervised training on surgical videos. Instead of designing a new pretext task, we provide a new insight that to transfer knowledge from large public dataset improves self-training on surgical data.
- We propose a semantic-preserving training to train a teacher model, which contains representation ability from ImageNet while producing accurate semantic information for the surgical data.
- We propose a distillation objective to enforce the similarity matrix between the teacher model and the student model to be consistent, which makes the student model to learn extra information during self-training.

2 Methodology

This section is divided into three main parts. Firstly, we introduce contrastive learning to perform self-supervised learning on surgical videos in Sect. 2.1. Then, the semantic-preserving training will be presented in Sect. 2.2. Finally, we will describe the distilled self-supervised learning, which transfers useful semantics from publicly available models to the self-training on surgical videos.

2.1 Contrastive Learning on Surgical Videos

Given a frame v sampled from the surgical videos, we apply two different transformations to it, which can be formally as:

$$x^q = t_q(v), \quad x^k = t_k(v), \quad (1)$$

where $t_q(\cdot)$ and $t_k(\cdot)$ are two transformations, sampled from the same transformation distribution. Then, we feed x^q into a query encoder $f_q(\cdot)$, followed by a projection head $f_h(\cdot)$ consists of 2-layer multilayer perceptron (MLP), to obtain the query representation $q = h_q(f_q(x^q))$. Similarly, we can obtain $k_+ = h_k(f_k(x^k))$, where k_+ indicates the positive embedding for q . Similar as [4, 10], we obtain a set of encoded samples $B = \{k_i\}_{i=1}^M$ that are the keys of a dictionary, where k_i is the i -th key in the dictionary, and M is the size of the dictionary. There is a single key in the dictionary, *i.e.*, k_+ , that matches q . In this paper, we use InfoNCE [18], which measures the similarity via dot production, and our contrastive objective is as follows:

$$\mathcal{L}_{con} = - \sum_{i=1}^N \log \frac{\exp(\text{sim}(q, k_+)/\tau)}{\sum_{i=1}^M \exp(\text{sim}(q, k_i)/\tau)}, \quad (2)$$




Fig. 2. Left: A semantic-preserving training is adopted to train a teacher model, *i.e.*, f_q^t , f_k^t , h_q^t and h_k^t , from the free available model. **Right:** We introduce a distilled self-supervised learning to distill the knowledge from the teacher model to the student model, *i.e.*, f_q^s , f_k^s , h_q^s and h_k^s .

where $\text{sim}(q, k_+)$ is the cosine similarity between q and k_+ , and τ is the temperature that controls the concentration level of the sample distribution [7]. In this paper, the parameters of the query encoder are updated by back-propagation, while the key encoder is momentum-updated encoder [4, 10]. Let define the parameters of $f_q(\cdot)$, $h_q(\cdot)$, $f_k(\cdot)$ and $h_k(\cdot)$ as $\theta_q^f(\cdot)$, $\theta_q^h(\cdot)$, $\theta_k^f(\cdot)$ and $\theta_k^h(\cdot)$ respectively. Then, the momentum-updating can be formulated as:

$$\theta_k^f \leftarrow m\theta_k^f + (1 - m)\theta_q^f, \tag{3}$$

$$\theta_k^h \leftarrow m\theta_k^h + (1 - m)\theta_q^h, \tag{4}$$

where $m \in [0, 1)$ is a momentum coefficient, and is set to 0.999 following [4].

2.2 Semantic-Preserving Training of Free Models

Our motivation is to leverage the semantic information from public free self-supervised model (*i.e.*, a teacher model) trained on large datasets (*e.g.*, ImageNet) to improve the self-supervised learning, (*i.e.*, a student model) on surgical videos. To this end, the teacher model should have two properties. (1) It should retain representation ability learned from ImageNet. Directly self-training on surgical video would make the model forget the learned knowledge from ImageNet. (2) It should produce accurate structured knowledge for surgical videos. Since there is a clear domain gap between ImageNet and surgical videos, transferring the noisy distilled information would degrade the performance.

To achieve these two properties, we propose the semantic-preserving training of free models on surgical video, which is shown in the left of Fig. 2. More specifically, we first initialize the parameters of $f_q(\cdot)$ and $f_k(\cdot)$ from the ImageNet pre-trained model, *i.e.*, using the weights of the free self-supervised model: Moco v2,

which can be formulated as $f_q(\cdot) \rightarrow f_q^t(\cdot)$ and $f_k(\cdot) \rightarrow f_k^t(\cdot)$. During semantic-preserving training, we fixed the parameters of the backbone, *i.e.*, $f_q^t(\cdot)$ and $f_k^t(\cdot)$, to maintain the prior learned semantics. Furthermore, to produce accurate predictions for latter distillation under the large domain gap, we conduct contrastive learning on surgical video to update the parameters of the two-layer MLP, *i.e.*, $h_q^t(\cdot)$ and $h_k^t(\cdot)$, via the contrastive objective. In this way, the trained model would learn to use the prior knowledge from ImageNet to discriminate the surgical frames, and the predictions from it would contain rich semantic information.

2.3 Distilled Self-supervised Learning

Unlike existing self-supervised learning approaches that focus on designing new pretext tasks, we aim to transfer the semantics knowledge from free public self-supervised models into the surgical video self-training. As shown in the right of Fig. 2, besides the standard contrastive learning, we also introduce a distillation objective to transfer the semantics from the teacher model (trained in Sect. 2.2) to the student model, *i.e.*, the model learned by self-supervised on surgical videos. The similarity matrix (*i.e.*, $\text{sim}(q, k_i)$) obtained in Eq. 2 measures the similarities between the query and its keys. The teacher model can generate around 82.3% Acc. for surgical phase recognition, showing that the similarity matrix contains useful semantics information that can benefit to surgical video understanding.

To leverage this information, our idea is to force the similarity matrix of the teacher and student model to be consistent; therefore, the teacher model can provide additional supervision for training the student model, leading to performance improvement. Specifically, we regard the similarity matrix of the teacher model as the soft targets to supervise the self-training of the student. To obtain the soft targets, we apply Softmax with the temperature scale τ to the similarity matrix of the teacher model. We find that the soft targets can be formulated as:

$$p_i^t = \frac{\exp(\text{sim}(q^t, k_i^t)/\tau)}{\sum_{i=j}^M \exp(\text{sim}(q^t, k_j^t)/\tau)}, \quad (5)$$

which is very similar to Eq. 2. Similarly, we can also obtain the similarity matrix of the student model, which can be defined as p_i^s . Finally, the distillation loss is computed by the KL-divergence loss to measure the distribution of p_i^t and p_i^s as follows:

$$\mathcal{L}_{dis} = \sum_{i=1}^M p_i^t \log\left(\frac{p_i^t}{p_i^s}\right), \quad (6)$$

where p_i^s is the similarity matrix of the student model, computed as in Eq. 5. Note that the distillation loss can be regarded as a soft version of the contrastive loss (defined in Eq. 2). This soft version provides much more information per training case than hard targets [12].

Combined with the distillation loss and the contrastive loss, the overall object of the distilled self-supervised can be formulated as:

$$\mathcal{L} = \mathcal{L}_{con} + \lambda \mathcal{L}_{dis}, \quad (7)$$

where λ is the trade-off hyper-parameter.

3 Experiments

3.1 Datasets

Cholec80. The Cholec80 dataset [24] contains 80 cholecystectomy surgeries videos. All videos are recorded as 25 fps, and the resolution of each frame is 1920×1080 or 854×480 . We sample the videos into 5 fps for reducing computation cost. Following the same setting in previous works, we split the datasets into 40 videos for training and rest 40 videos for testing [14, 28].

MI2CAI16. There are 41 videos in the MI2CAI16 [21] dataset, and the videos in it are recorded at 25 fps. Among them, 27 videos are used for training and 14 videos are used for test, following previous approaches [14, 28].

3.2 Implementation Details

We use the ResNet50 [11] as the backbone to extract features for each frame. After that, following [6, 24, 28], a multi-stage temporal convolution (MS-TCN) [8] is used to extracted temporal relations for frame features and predicts phase results. Following the same evaluation protocols [13–15, 24], we employ four commonly-used metrics, i.e., accuracy (AC), precision (PR), recall (RE), and Jaccard (JA) to evaluate the phase prediction accuracy. To evaluate the self-supervised training performance, we self-train the ResNet50 backbone on surgical videos on training set of Cholec80 by different self-supervised learning approaches [2, 4, 5, 27]. Then, we linearly fine-tune the self-trained model on the training set of Cholec80 and MI2CAI16. After that, we extract features for each frame of videos by the fine-tuned model. Finally, based on the fixed extracted features, we use MS-TCN to predict results for surgical phase recognition. We set τ to 0.07 and λ to 5 respectively.

3.3 Comparisons with the State-of-the-Art Methods

To prove the effect of our proposed method, we compare our method with the state-of-the-art self-supervised learning approaches, *e.g.*, MoCo v2 [4], SwAV [2], SimSiam [5] and RegionCL [27]. The results are shown in Table 1. Note that, our method uses MoCo v2 [4] as the self-supervised learning baseline. We find that self-supervised training on ImageNet outperforms the model trained from scratch with a clear margin. This indicates the powerful representation of the ImageNet pre-training, and motivates us to leverage semantics from it to improve the self-supervised training on surgical data. It can be also found that self-supervised

Table 1. Results on Cholec80 dataset and M2CAI16 dataset. ‘‘Scratch’’ indicates training from scratch. ‘‘IN-Sup’’ and ‘‘IN-MoCo v2’’ refer to supervised learning and self-supervised learning on ImageNet, respectively.

Methods	Cholec80				M2CAI16			
	Accuracy	Precision	Recall	Jaccard	Accuracy	Precision	Recall	Jaccard
Scratch	65.0 \pm 13.8	67.7 \pm 14.3	53.0 \pm 21.4	40.7 \pm 18.4	59.8 \pm 12.9	65.0 \pm 13.6	56.6 \pm 29.7	40.0 \pm 23.0
Models Trained on ImageNet								
IN-Sup	83.1 \pm 10.5	64.7 \pm 11.0	82.3 \pm 7.6	77.8 \pm 9.8	76.0 \pm 6.3	78.8 \pm 5.2	73.8 \pm 17.6	58.9 \pm 15.0
IN-MoCov2	82.3 \pm 9.8	64.2 \pm 12.0	81.2 \pm 6.1	77.7 \pm 10.5	75.1 \pm 5.8	77.8 \pm 6.4	73.5 \pm 10.4	58.2 \pm 14.2
Models Trained on Surgical Videos								
MoCo v2 [4]	85.7 \pm 10.4	72.0 \pm 9.1	85.4 \pm 6.3	84.9 \pm 5.4	78.4 \pm 4.6	80.5 \pm 5.1	78.5 \pm 12.6	63.5 \pm 13.6
SwAV [2]	84.7 \pm 2.8	71.6 \pm 7.9	84.9 \pm 7.6	84.2 \pm 6.3	77.3 \pm 6.1	79.9 \pm 6.3	77.8 \pm 11.3	63.1 \pm 11.5
SimSiam [5]	83.9 \pm 6.7	68.5 \pm 5.5	84.5 \pm 5.0	82.6 \pm 5.7	77.0 \pm 5.3	79.1 \pm 6.2	77.1 \pm 8.2	62.8 \pm 9.7
RegionCL [27]	84.8 \pm 11.4	69.8 \pm 11.1	85.0 \pm 6.5	82.8 \pm 10.9	78.2 \pm 10.4	79.9 \pm 8.1	78.0 \pm 9.5	63.4 \pm 7.4
Ours	87.3 \pm 9.5	73.5 \pm 9.9	85.1 \pm 7.6	86.1 \pm 6.9	81.1 \pm 6.5	80.7 \pm 8.3	83.2 \pm 6.9	67.9 \pm 5.1

Table 2. Ablation study on different transferring methods on Cholec80 dataset.

Method	Accuracy	Precision	Recall	Jaccard
Addition	84.0 \pm 10.5	67.3 \pm 11.1	82.7 \pm 8.7	81.6 \pm 8.3
Concatenation	84.2 \pm 9.4	67.9 \pm 11.7	83.1 \pm 8.6	81.0 \pm 10.1
Initialization	85.5 \pm 10.9	71.7 \pm 8.1	83.8 \pm 7.2	83.2 \pm 8.4
Distillation	87.3 \pm 9.5	73.5 \pm 9.9	85.1 \pm 7.6	86.1 \pm 6.9

training on surgical videos outperforms that training on ImageNet [20]. Combined with the distillation of semantics from ImageNet, our method can improve the baseline, *i.e.*, MoCo v2, over 1.6% on Cholec80.

3.4 Ablation Studies

Effect of Distillation. Table 2 compares the performance of different transferring methods on Cholec80. Specifically, we denote the feature maps generated from the teacher model and our model as \mathbf{F}^t and \mathbf{F}^s respectively. ‘Addition’ means the sum of \mathbf{F}^t and \mathbf{F}^s , *i.e.*, $\mathbf{F}^t + \mathbf{F}^s$. Similarly, ‘Concatenation’ indicates the concatenation of \mathbf{F}^t and \mathbf{F}^s . ‘Initialization’ means our model is initialized by the teacher model. ‘Distillation’ means using the distillation objective defined in Eq. 6. It is clear that the distillation model is the best way to transfer the knowledge from the teacher model.

Effect of the Semantic-Preserving Training of Free Models. We conduct an ablation study on Cholec80 to evaluate the effect of the semantic-preserving training, and report the results in Table 3. It is clear that not training the teacher model, *i.e.*, the first row in Table 3, cannot predict correct knowledge for surgical data, and would degrade the performance, *e.g.*, only achieve 84.7% accuracy. Furthermore, training too much would hurt the semantics learned from ImageNet; See the third and the last rows in Table 3. We can find that our proposed semantic-preserving training (only updating parameters of the head projection




Fig. 3. Accuracy for surgical phase recognition for different methods and varied sizes of label fractions on **Left:** Cholec80 and **Right:** MI2CAI16. “Scratch” indicates training from scratch. “Mo-IN” and “Mo-SU” refer to training from the model pre-trained on ImageNet and surgical data by MoCo v2, respectively.

Table 3. Ablation study on the semantic-preserving training on Cholec80 dataset. ‘Backbone’ indicates training the backbone of the model, and ‘Head’ means training the head projection of the model.

Backbone	Head	Accuracy	Precision	Recall	Jaccard
\times	\times	84.7 \pm 11.0	67.5 \pm 11.4	84.2 \pm 8.2	79.6 \pm 8.8
\times	\checkmark	87.3 \pm 9.5	73.5 \pm 9.9	85.1 \pm 7.6	86.1 \pm 6.9
\checkmark	\times	86.1 \pm 11.8	72.4 \pm 9.4	86.1 \pm 6.5	84.8 \pm 9.3
\checkmark	\checkmark	85.2 \pm 10.9	71.3 \pm 10.4	85.1 \pm 5.7	85.8 \pm 6.5

(*i.e.*, h_q^t and h_k^t) and fixing the backbone (*i.e.*, f_q^t and f_k^t) can achieve the best performance. Since trained with this way, the teacher model would not only maintain the semantics learned from ImageNet, but also can predict accurate knowledge for transferring.

Label-Efficiency Analysis. To evaluate the label-efficiency of our self-training model, we fine-tune the pre-trained model on different fractions of labeled training data. Note that in this ablation study, we train all parameters of the pre-trained model. As shown in Fig. 3, our proposed approach achieves the best label-efficiency performance. Furthermore, fine-tuning with fewer labels, the improvement gain is larger.

4 Conclusion

Unlike existing self-supervised papers that design new pretext tasks, this paper proposes a novel free lunch method for self-supervised learning with surgical videos by distilling knowledge from free available models. Our motivation is that the publicly released model trained on natural images shows comparable performance with the model trained on the surgical data. To leverage the rich semantics from the available models, we propose a semantic-persevering method

to train a teacher model that contains prior information and can produce accurate predictions for surgical data. Finally, we introduce a distillation objective to the self-training on surgical data to enable the model to learn extra knowledge from the teacher model. Experimental results indicate that our method can improve the performance of the existing self-supervised learning methods.

Acknowledgement. This work was supported by a research grant from HKUST-BICI Exploratory Fund (HCIC-004) and a research grant from Shenzhen Municipal Central Government Guides Local Science and Technology Development Special Funded Projects (2021Szvup139), and under the RIE2020 Industry Alignment Fund - Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

1. Blum, T., Feußner, H., Navab, N.: Modeling and segmentation of surgical workflow from laparoscopic video. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010. LNCS, vol. 6363, pp. 400–407. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15711-0_50
2. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. *Adv. Neural. Inf. Process. Syst.* **33**, 9912–9924 (2020)
3. Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., Rueckert, D.: Self-supervised learning for medical image analysis using image context restoration. *Med. Image Anal.* **58**, 101539 (2019)
4. Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. *arXiv preprint [arXiv:2003.04297](https://arxiv.org/abs/2003.04297)* (2020)
5. Chen, X., He, K.: Exploring simple siamese representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15750–15758 (2021)
6. Ding, X., Li, X.: Exploiting segment-level semantics for online phase recognition from surgical videos. *arXiv preprint [arXiv:2111.11044](https://arxiv.org/abs/2111.11044)* (2021)
7. Ding, X., et al.: Support-set based cross-supervision for video grounding. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 11573–11582 (2021)
8. Farha, Y.A., Gall, J.: MS-TCN: multi-stage temporal convolutional network for action segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3575–3584 (2019)
9. Gao, X., Jin, Y., Long, Y., Dou, Q., Heng, P.A.: Trans-SVNet: accurate phase recognition from surgical videos via hybrid embedding aggregation transformer. *arXiv preprint [arXiv:2103.09712](https://arxiv.org/abs/2103.09712)* (2021)
10. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738 (2020)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778 (2016)
12. Hinton, G., Vinyals, O., Dean, J., et al.: Distilling the knowledge in a neural network. *arXiv preprint [arXiv:1503.02531](https://arxiv.org/abs/1503.02531)*, vol. 2, no. 7 (2015)

13. Jin, Y., et al.: SV-RCNet: workflow recognition from surgical videos using recurrent convolutional network. *IEEE Trans. Med. Imaging* **37**(5), 1114–1126 (2017)
14. Jin, Y., et al.: Multi-task recurrent convolutional network with correlation loss for surgical video analysis. *Med. Image Anal.* **59**, 101572 (2020)
15. Jin, Y., Long, Y., Chen, C., Zhao, Z., Dou, Q., Heng, P.A.: Temporal memory relation network for workflow recognition from surgical video. *IEEE Trans. Med. Imaging* **40**(7), 1911–1923 (2021)
16. Li, X., et al.: Rotation-oriented collaborative self-supervised learning for retinal disease diagnosis. *IEEE Trans. Med. Imaging* **40**(9), 2284–2294 (2021)
17. Li, X., Jia, M., Islam, M.T., Yu, L., Xing, L.: Self-supervised feature learning via exploiting multi-modal data for retinal disease diagnosis. *IEEE Trans. Med. Imaging* **39**(12), 4023–4033 (2020)
18. Van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv e-prints* pp. arXiv-1807 (2018)
19. Pan, T., Song, Y., Yang, T., Jiang, W., Liu, W.: Videomoco: contrastive video representation learning with temporally adversarial examples. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11205–11214 (2021)
20. Russakovsky, O., et al.: Imagenet large scale visual recognition challenge. *Int. J. Comput. Vision* **115**(3), 211–252 (2015)
21. Stauder, R., Ostler, D., Kranzfelder, M., Koller, S., Feußner, H., Navab, N.: The TUM LapChole dataset for the M2CAI 2016 workflow challenge. *arXiv preprint arXiv:1610.09278* (2016)
22. Taleb, A., et al.: 3D self-supervised methods for medical imaging. *Adv. Neural. Inf. Process. Syst.* **33**, 18158–18172 (2020)
23. Tian, Y., Krishnan, D., Isola, P.: Contrastive multiview coding. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) *ECCV 2020*. LNCS, vol. 12356, pp. 776–794. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58621-8_45
24. Twinanda, A.P., Shehata, S., Mutter, D., Marescaux, J., De Mathelin, M., Padoy, N.: EndoNet: a deep architecture for recognition tasks on laparoscopic videos. *IEEE Trans. Med. Imaging* **36**(1), 86–97 (2016)
25. Wang, Z., Ding, X., Zhao, W., Li, X.: Less is more: surgical phase recognition from timestamp supervision. *arXiv preprint arXiv:2202.08199* (2022)
26. Xie, X., Chen, J., Li, Y., Shen, L., Ma, K., Zheng, Y.: Instance-aware self-supervised learning for nuclei segmentation. In: Martel, A.L., et al. (eds.) *MICCAI 2020*. LNCS, vol. 12265, pp. 341–350. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59722-1_33
27. Xu, Y., Zhang, Q., Zhang, J., Tao, D.: RegionCL: can simple region swapping contribute to contrastive learning? *arXiv preprint arXiv:2111.12309* (2021)
28. Yi, F., Jiang, T.: Not end-to-end: explore multi-stage architecture for online surgical phase recognition. *arXiv preprint arXiv:2107.04810* (2021)
29. Zappella, L., Béjar, B., Hager, G., Vidal, R.: Surgical gesture classification from video and kinematic data. *Med. Image Anal.* **17**(7), 732–745 (2013)
30. Zhou, H.Y., Lu, C., Yang, S., Han, X., Yu, Y.: Preservational learning improves self-supervised medical image models by reconstructing diverse contexts. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3499–3509 (2021)
31. Zhou, H.-Y., Yu, S., Bian, C., Hu, Y., Ma, K., Zheng, Y.: Comparing to learn: surpassing imagenet pretraining on radiographs by comparing image representations. In: Martel, A.L., et al. (eds.) *MICCAI 2020*. LNCS, vol. 12261, pp. 398–407. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_39

32. Zhou, Z., Sodha, V., Pang, J., Gotway, M.B., Liang, J.: Models genesis. *Med. Image Anal.* **67**, 101840 (2021)
33. Zhu, J., Li, Y., Hu, Y., Ma, K., Zhou, S.K., Zheng, Y.: Rubik's Cube+: a self-supervised feature learning framework for 3D medical image analysis. *Med. Image Anal.* **64**, 101746 (2020)
34. Zhuang, X., Li, Y., Hu, Y., Ma, K., Yang, Y., Zheng, Y.: Self-supervised feature learning for 3D medical images by playing a Rubik's cube. In: Shen, D., et al. (eds.) *MICCAI 2019. LNCS*, vol. 11767, pp. 420–428. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32251-9_46