



# Imbalance-Aware Self-supervised Learning for 3D Radiomic Representations

Hongwei Li<sup>1,2</sup>, Fei-Fei Xue<sup>4</sup>, Krishna Chaitanya<sup>3</sup>, Shengda Luo<sup>5</sup>, Ivan Ezhov<sup>1</sup>, Benedikt Wiestler<sup>6</sup>, Jianguo Zhang<sup>4</sup>(✉), and Bjoern Menze<sup>1,2</sup>

<sup>1</sup> Department of Computer Science, Technical University of Munich, Munich, Germany  
hongwei.li@tum.de

<sup>2</sup> Department of Quantitative Biomedicine, University of Zurich, Zürich, Switzerland  
<sup>3</sup> ETH Zurich, Zürich, Switzerland

<sup>4</sup> Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen, China  
zhangjg@sustech.edu.cn

<sup>5</sup> Faculty of Information Technology, Macau University of Science and Technology, Macao, China

<sup>6</sup> Klinikum rechts der Isar, Technical University of Munich, Munich, Germany

**Abstract.** Radiomics can quantify the properties of regions of interest in medical image data. Classically, they account for pre-defined statistics of shape, texture, and other low-level image features. Alternatively, deep learning-based representations are derived from supervised learning but require expensive annotations and often suffer from overfitting and data imbalance issues. In this work, we address the challenge of learning the representation of a 3D medical image for an effective quantification under data imbalance. We propose a *self-supervised* representation learning framework to learn high-level features of 3D volumes as a complement to existing radiomics features. Specifically, we demonstrate how to learn image representations in a self-supervised fashion using a 3D Siamese network. More importantly, we deal with data imbalance by exploiting two unsupervised strategies: a) sample re-weighting, and b) balancing the composition of training batches. When combining the learned self-supervised feature with traditional radiomics, we show significant improvement in brain tumor classification and lung cancer staging tasks covering MRI and CT imaging modalities. Codes are available in <https://github.com/hongweilibran/imbalanced-SSL>.

---

**Electronic supplementary material** The online version of this chapter ([https://doi.org/10.1007/978-3-030-87196-3\\_4](https://doi.org/10.1007/978-3-030-87196-3_4)) contains supplementary material, which is available to authorized users.

## 1 Introduction

Great advances have been achieved in supervised deep learning, reaching expert-level performance on some considerably challenging applications [11]. However, supervised methods for image classification commonly require relatively large-scale datasets with ground-truth labels which is time- and resource-consuming in the medical field. Radiomics is a translational field aiming to extract objective and quantitative information from clinical imaging data. While traditional radiomics methods, that rely on statistics of shape, texture and others [1], are proven to be generalizable in various tasks and domains, their discriminativeness is often not guaranteed since they are low-level features which are not specifically optimized on target datasets.

Self-supervised learning for performing *pre-text tasks* have been explored in medical imaging [24, 25], that serve as a proxy task to pre-train the deep neural networks. They learn representations commonly in a supervised manner on proxy tasks. Such methods depend on heuristics to design pre-text tasks which could limit the discriminativeness of the learnt representations. In this work, we investigate self-supervised *representation learning* which aims to **directly** learn the representation of the data without a proxy task.

Recent contrastive learning-based methods [6, 15, 20] learn informative representations *without* human supervision. However, they often rely on large batches to train and most of them work for 2D images. To this end, due to the high dimensionality and limited number of training samples in medical field, applying contrastive learning-based methods may not be practically feasible in 3D datasets. Specially, in this study, we identify two main differences required to adapt self-supervised representation learning for radiomics compared to natural image domain: i) Medical datasets are often multi-modal and three dimensional. Thus, learning representation methods in 3D medical imaging would be computationally expensive. ii) heterogeneous medical datasets are *inherently imbalanced*, e.g. distribution disparity of disease phenotypes. Existing methods are built upon approximately balanced datasets (e.g. CIFAR [18] and ImageNet [10]) and do not assume the existence of data imbalance. Thus, how to handle data imbalance problem is yet less explored in the context of self-supervised learning.

**Related Work.** Radiomic features have drawn considerable attention due to its predictive power for treatment outcomes and cancer genetics in personalized medicine [12, 23]. Traditional radiomics include shape features, first-, second-, and higher- order statistics features.

Self-supervised representation learning [3, 6, 7, 13, 15, 22] have shown steady improvements with impressive results on multiple natural image tasks, mostly based on contrastive learning [14]. Contrastive learning aims to attract positive (or *similar*) sample pairs and rebuff negative (or *disimilar*) sample pairs. Positive sample pairs can be obtained by generating two augmented views of one sample, and the remaining samples in the batch can be used to construct the negative samples/pairs for a given positive pair. In practice, contrastive learning

methods benefit from a large number of negative samples. In medical imaging, there are some existing work related to contrastive learning [2, 17, 21]. Chaitanya *et al.*'s work [5] is the most relevant to our study, which proposed a representation learning framework for image segmentation by exploring local and global similarities and dissimilarities. Though these methods are effective in learning representations, they require a large batch size and/or negative pairs, which make them difficult to apply to 3D medical data. Chen *et al.* [8] demonstrates that a Siamese network can avoid the above issues on a 2D network. The Siamese network, which contains two encoders with shared weights, compares two similar representations of two augmented samples from one sample. Importantly, it neither uses negative pairs nor a large batch size. Considering these benefits, we borrow the Siamese structure and extend it to 3D *imbalanced* medical datasets.

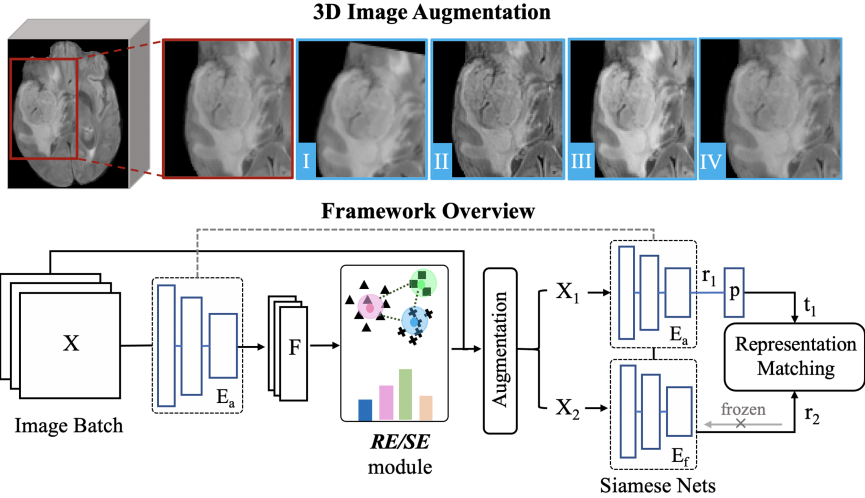
**Contributions.** Our contribution is threefold: (1) We develop a 3D Siamese network to learn self-supervised representation which is high-level and discriminative. (2) For the first time, we explore how to tackle the data imbalance problem in self-supervised learning without using labels and propose two effective unsupervised strategies. (3) We demonstrate that self-supervised representations can complement the existing radiomics and the combination of them outperforms supervised learning in two applications.

## 2 Methodology

The problem of interest is how to learn high-level, discriminative representations on 3D imbalanced medical image datasets in a self-supervised manner. The schematic view of the framework is illustrated in Fig. 1. First, a pre-trained 3D encoder network, denoted as  $E_a$ , takes a batch of original images  $X$  with batch size  $N$  as input and outputs  $N$  representation vectors. The details of the 3D encoder is shown in Table 4 of the Supplementary. The features are fed into the *RE/SE* module to estimate their individual weights or to resample the batch.

Then each image  $x$  in the batch  $X$  is randomly augmented into two images (or called an image pair). They are processed and compared by a 3D Siamese network, nicknamed *3DSiam*, which enjoys relatively low memory without relying on large training batch of 3D data. The proposed *3DSiam* extends original 2D Siamese network [8] from processing 2D images to 3D volumes while inherits its advantages. Since medical datasets are inherently imbalanced, by intuition sole *3DSiam* would probably suffer from imbalanced data distribution. In the following, we first introduce *RE/SE* module to mitigate this issue.

**RE/SE Module to Handle Imbalance.** Since there is no prior knowledge on the data distribution available, the way to handle imbalance must be *unsupervised*. The vectors mentioned above are fed into a *RE/SE* module before training the *3DSiam* network. The  $k$ -means algorithm is used first to cluster the representation vectors into  $k$  centers. We then proposed two simple yet effective strategies: a) sample re-weighting (RE), and b) sample selection (SE):



**Fig. 1.** Our proposed framework learns invariance from extensive 3D image augmentation within four categories: I) affine transform, II) appearance enhancement, III) contrast change, and IV) adding random noise. First, an image batch  $X$  is fed into an initialized 3D encoder to obtain its representation  $F$ . The RE/SE module first estimates its distribution by k-means based clustering and uses two strategies including sample re-weighting (RE) or sample selection (SE) to alleviate data imbalance issue. Each image is randomly augmented into two positive samples  $\{X_1, X_2\}$  which are then used to train a 3D Siamese network by comparing their representations from two encoders  $\{E_a, E_f\}$  with shared weights.  $p$  is a two-layer perceptron to transform the feature.

- a) Sample re-weighting (RE). Denote a batch with  $N$  samples as  $X = \{x_i | i = 1, 2, \dots, N\}$ . Given  $k$  clusters, denote the distribution of  $k$  clusters of features as  $F = \{f_j | j = 1, 2, \dots, k\}$  over  $N$  samples.  $f_j$  denotes the frequency of cluster  $j$ . Then we assign different weights to the samples in each cluster  $j$ . For each sample  $x_i$ , representation vector of which belongs to cluster  $j$ , we assign it a weight of  $N/f_j$  to penalize the imbalanced distribution during the batch training. In practice, we further normalize it by re-scaling to guarantee the minimum weight is 1 for each input batch.
- b) Sample selection (SE). Denoting the clusters' centroids as  $C = \{c_1, c_2, \dots, c_k\}$ , we find the maximum Euclidean distance  $\max_{i,j \in [1,k], i \neq j} d(c_i, c_j)$  among all pairs of centroids.  $k$  is a hyper-parameter here. We hypothesize that the clusters with maximum centroid distance are representation vectors from different groups. To select  $m$  samples from the original  $N$  samples to form a new batch, denoted by  $B_c = \{x_1, x_2, \dots, x_m\}$ , we sample  $\frac{m}{2}$  nearest sample points centered on each of the selected maximum-distance centroids.  $m$  is set to be smaller than  $\frac{N}{k}$  for low computation complexity and efficient sampling. The selected new batch is then used to train our  $3DSiam$  network. A motivation behind the selection strategy is outlined in Supplementary.

**3D Siamese Network.** The *3DSiam* takes as input two randomly augmented views  $x_1$  and  $x_2$  from a sample  $x$ . The two views are processed by two 3D encoder networks with *shared* weights. One of the encoder has frozen weights when training (denoted as  $E_f$ ) and the other one is with active weights (denoted as  $E_a$ ). Before training,  $E_f$  is always updated to the weights of  $E_a$ .  $E_a$  is followed by a two-layer perceptron called *predictor* to transform the features. The final objective is to optimize a matching score between the two *similar* representations  $t_1 \triangleq p(E_a(x_1))$  and  $r_2 \triangleq E_f(x_2)$ . *3DSiam* minimizes their negative cosine similarity, which is formulated as:

$$S(t_1, r_2) = -\frac{t_1}{\|t_1\|_2} \cdot \frac{r_2}{\|r_2\|_2}, \quad (1)$$

where  $\|\cdot\|_2$  is  $L_2$ -norm. Following [6], we define a symmetrized loss to train the two encoders, formulated as:

$$\mathcal{L} = \frac{1}{2}S(t_1, r_2) + \frac{1}{2}S(t_2, r_1), \quad (2)$$

where  $t_2 \triangleq p(E_f(x_2))$ ,  $r_1 \triangleq E_f(x_1)$ . This loss is defined and computed for each sample with re-weighting in the batch  $X$  or the new batch  $B_c$  with equal weights. Notably the encoder  $E_f$  on  $x_2$  receives no gradient from  $r_2$  in the first term of Eq. (2), but it receives gradients from  $t_2$  in the second term (and vice versa for  $x_1$ ). This training strategy avoids collapsing solutions, i.e.,  $t_1$  and  $r_2$  outputs a constant over the training process. When training is finished,  $r_2$  is used as the final representation.

### 3 Experiments

**Datasets and Preprocessing.** The evaluation of our approach is performed on two public datasets: 1) a multi-center MRI dataset (*BraTS*) [4, 19] including 326 patients with brain tumor. The MRI modalities include FLAIR, T1, T2 and T1-c with a uniform voxel size  $1 \times 1 \times 1 \text{ mm}^3$ . Only FLAIR is used in our experiment for simplicity of comparisons. 2) a lung CT dataset with 420 non-small cell lung cancer patients (*NSCLC-radiomics*) [1, 9]<sup>1</sup>. The effectiveness of the learnt representations is evaluated on two classification tasks (also called ‘down-stream task’): a) discriminating high grade (H-grade) and low grade tumor (L-grade), and b) predicting lung cancer stages (i.e. I, II or III). The *BraTS* dataset is imbalanced in two aspects: a) the distribution of ground truth labels (H-grade *vs.* L-grade); b) the distribution of available scans among different medical centers. For *NSCLC-radiomics*, the distribution of ground truth labels are imbalanced as well, with ratio of 2:1:6 for stage I, II and III respectively. For *BraTS*, we make use of the segmentation mask to get the centroid and generate a 3D bounding box of  $96 \times 96 \times 96$  to localize the tumour. If the bounding box exceeds the original volume, the out-of-box region was padded with background

<sup>1</sup> Two patients were excluded as the ground truth labels are not available.

intensity. For *NSCLC-radiomics*, we get the lung mask by a recent public lung segmentation tool [16], and then generate a  $224 \times 224 \times 224$  bounding box to localize the lung. The lung volume was then resized to  $112 \times 112 \times 112$  due to memory constraint. The intensity range of all image volumes was rescaled to  $[0, 255]$  to guarantee the success of intensity-based image transformations.

**Configuration of the Training Schedule.** We build a 3D convolutional neural network with two bottleneck blocks as the encoder for all experiments (details in Supplementary). In the beginning, we pre-train *3DSiam* for one epoch with a batch size of 6. Then we use it to extract features for the *RE/SE* module. After first epoch, the encoder from the *last iteration* is employed for *dynamic* feature extraction. For down-stream task evaluations, we use the last-layer feature of the encoder. For 3D data augmentation, we apply four categories shown in Fig. 1, including random rotations in  $[-20, 20]$  degrees, random scale between  $[0.7, 1.3]$ , and random shift between  $[-0.05, 0.05]$ , Gamma contrast adjustment between  $[0.7, 1.5]$ , image sharpening, and Gaussian blurring, considering the special trait of medical images. For optimization, we use Adam with  $10^{-2}$  learning rate and  $10^{-4}$  weight decay. Each experiment is conducted using one Nvidia RTX 6000 GPU with 24 GB memory. The number of cluster  $k$  is set to 3 in all experiments. Its effect is analyzed in the last section.

*Computation Complexity.* For *3DSiam* without *RE/SE* module, the training takes only around four hours for 50 epochs for the results reported for brain tumor classification task. We do not observe significant improvement when increasing the number of epochs after 50. We train *3DSiam* with *RE/SE* module for around 2000 iterations (not epochs) to guarantee that similar number of training images for the models of comparison are involved. In *RE/SE* module, the main computation cost is from  $k$ -means algorithm. We have observed that the overall computation time has increased by 20% (with i5-5800K CPU). It is worth noting that *RE/SE* module is not required during the inference stage, thus there is no increase of the computational cost in testing.

**Feature Extraction and Aggregation.** For each volume, we extract a set of 107 traditional radiomics features<sup>2</sup> including first- and second-order statistics, shape-based features and gray level co-occurrence matrix, denoted as  $f_{trad}$ . For the self-supervised learning one, we extract 256 features from the last fully connected layer of the encoder  $E_a$ , denoted as  $f_{SSL}$ . To directly evaluate the effectiveness of SSL-based features, we concatenate them to a new feature vector  $f = [f_{trad}, f_{SSL}]$ . Note that  $f_{trad}$  and  $f_{SSL}$  are always from the same subjects.

**Evaluation Protocol, Classifier and Metrics.** For evaluation, we follow the common protocol to evaluate the quality of the pre-trained representations by training a *supervised* linear support vector machine (SVM) classifier on the

<sup>2</sup> <https://github.com/Radiomics/pyradiomics>.

training set, and then evaluating it on the test set. For binary classification task (BraTS), we use the sensitivity and specificity as the evaluation metrics. For multi-class classification task (lung cancer staging), we report the overall accuracy and minor-class (i.e. stage II) accuracy considering all testing samples. We use *stratified* five-fold cross validation to reduce selection bias and validate each model. In each fold, we randomly sample 80% subjects from each class as the training set, and the remaining 20% for each class as the test set. Within each fold, we employ 20% of training data to optimize the hyper-parameters.

## 4 Results

**Quantitative Comparison.** We evaluate the effectiveness of the proposed self-supervised radiomics features on two classification tasks: a) discrimination of low grade and high grade of brain tumor and b) staging of lung cancer.

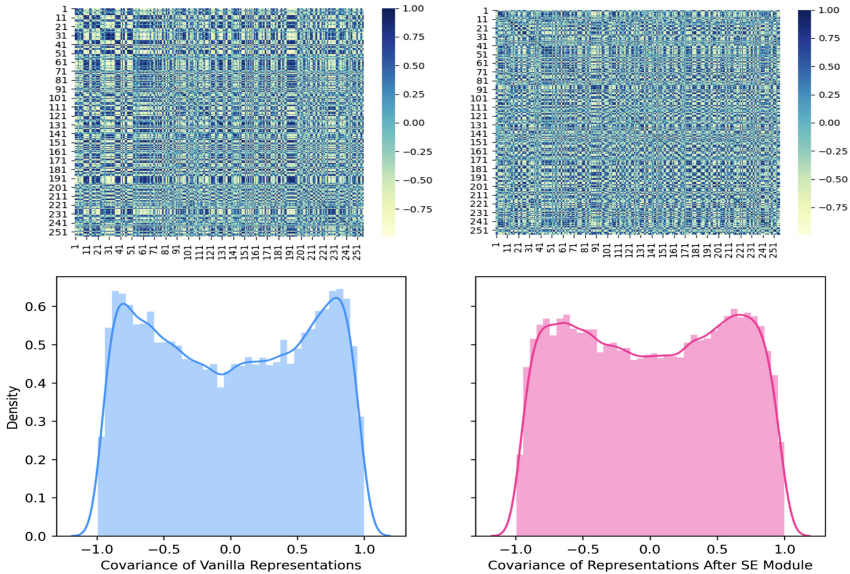
**Table 1.** Comparison of the performances of different kinds of features in two downstream tasks using stratified cross-validation. We further reduce 50% training data in each fold of validation to show the effectiveness against supervised learning. Our method outperforms supervised learning in both scenarios.

| Methods             | <i>BraTS</i>            |                     | <i>Lung cancer staging</i>   |                     |
|---------------------|-------------------------|---------------------|------------------------------|---------------------|
|                     | Sensitivity/Specificity |                     | Overall/Minor-class accuracy |                     |
|                     | <i>Full labels</i>      | <i>50% labels</i>   | <i>Full labels</i>           | <i>50% labels</i>   |
| Trad. radiomics     | 0.888/0.697             | <b>0.848</b> /0.697 | 0.490/0.375                  | 0.481/0.325         |
| Rubik’s cube [25]   | 0.744/0.526             | 0.680/0.486         | 0.459/0.325                  | 0.433/0.275         |
| 3DSiam              | 0.844/0.407             | 0.808/0.526         | 0.459/0.300                  | 0.445/0.300         |
| 3DSiam+SE           | 0.848/0.513             | 0.824/0.566         | 0.471/0.350                  | 0.443/0.325         |
| 3DSiam+RE           | 0.868/0.486             | 0.828/0.605         | 0.459/0.375                  | 0.445/0.325         |
| Trad.+3DSiam        | 0.904/0.645             | 0.804/0.566         | 0.495/0.350                  | 0.486/0.350         |
| Trad.+3DSiam+SE     | 0.916/ <b>0.711</b>     | <b>0.848/0.763</b>  | <b>0.538</b> /0.375          | <b>0.519</b> /0.350 |
| Trad.+3DSiam+RE     | <b>0.920/0.711</b>      | 0.804/0.763         | 0.524/ <b>0.425</b>          | 0.502/ <b>0.40</b>  |
| Supervised learning | 0.888/0.711             | 0.804/0.566         | 0.526/0.375                  | 0.467/0.325         |

*Effectiveness of RE/SE Module.* From the first row of Table 1, one can observe that traditional radiomics itself brings powerful features to quantify tumor characteristics. On BraTS dataset, the comparison between traditional radiomics and vanilla self-supervised radiomics (*3DSiam*) confirms our hypothesis that features learned by vanilla self-supervised method behave poorly, especially on the minor class (poor specificity). However, self-supervised radiomics with *RE* or *SE* module surpasses *3DSiam* in specificity by a large margin. The aggregation of the vanilla self-supervised representation and traditional radiomics does not show significant improvement. More importantly, with *RE/SE* module added, the specificity increased by 6.6%, from 64.5% to 71.1%, which indicates a large

boost in predicting the minor class (i.e. L-grade). Both comparisons of rows 4, 5, 6 and rows 7, 8, 9 demonstrate the success of our RE/SE module in tackling class imbalance, i.e., promoting the recognition of the minor class, while preserving the accuracy of the major class.

*Comparison with State-of-the-Art.* Our method (*Trad.+3DSiam+SE* in Table 2) outperforms the supervised one in two scenarios in two classification tasks, the result of which is achieved by using the same encoder backbone with a weighted cross-entropy loss. When it is trained with 50% less labels, the performance of supervised model decrease drastically. On lung cancer staging with three classes, although the overall accuracy of self-supervised radiomics is lower than the traditional one, with the *RE/SE* module, the combination of two kinds of radiomics achieves the topmost overall accuracy. This demonstrates the proposed self-supervised radiomics is complementary to existing radiomics. In the second row, we show the result of one self-supervised learning method trained by playing Rubik cubes [25] to learn contextual information with a same encoder. We observe that the representation learned in proxy task is less discriminative than the one directly from representation learning.



**Fig. 2.** Covariance analysis of the representations before and after the SE module. Across all 326 tumor patients, each feature was correlated with other ones, thereby generating the correlation coefficients. The density map show that the vanilla representation before SE module are more correlated (redundant) than the one after.



## Analysis of Representations and Hyperparameters

*Feature Covariance.* For a better understanding of the role of the proposed module in relieving data imbalance problem, we further analyze the feature covariance to understand the role of the *SE module*. Consider two paired variables  $(x_i, x_j)$  in the representation  $R$ . Given  $n$  samples  $\{(x_{i1}, x_{j1}), (x_{i2}, x_{j2}), \dots, (x_{in}, x_{jn})\}$ , Pearson’s correlation coefficient  $r_{x_i x_j}$  is defined as:  $r_{x_i x_j} = \frac{cov(x_i, x_j)}{\sigma_{x_i} \sigma_{x_j}}$ , where  $cov$  is the covariance and  $\sigma$  is the standard deviation. We found that the features after *SE* module become more compact as shown in Fig. 2 and more discriminative compared to the features without *SE* module.

*Effect of the Number of Clusters  $k$ .* The hyper-parameter  $k$  in the *SE* module is the number of clusters, which plays a vital role in constructing new batch. To evaluate its effect, we use different  $k$  to train *3DSiam* and evaluate it through classification task. To fairly compare different values of  $k$ , we keep the size  $m$  of the new batch  $B_c$  fixed to 6 which is also the batch size when  $k = 0$  (without *SE* module). The initial batch size  $N$  is set to  $k \times q$  where  $q$  is empirically set to 10 in the comparison. The AUC achieves the highest when  $k = 3$ . With  $k = 5$ , the AUC drops. This is probably because when  $k$  becomes large, the sampling may be biased when only considering a pair of clustering centers. For details, please refer to the curves of AUC over the number of clusters in Table 3 in Supplementary.

## 5 Conclusion

In this work, we proposed a 3D *self-supervised* representation framework for medical image analysis. It allows us to learn effective 3D representations in a self-supervised manner while considering the imbalanced nature of medical datasets. We have demonstrated that data-driven self-supervised representation could enhance the predictive power of radiomics learned from *large-scale* datasets without annotations and could serve as an effective compliment to the existing radiomics features for medical image analysis. Dealing with imbalance is an important topic and we will explore other strategies in the future.

**Acknowledgement.** This work was supported by Helmut Horten Foundation. I. E. was supported by the TRABIT network under the EU Marie Skłodowska-Curie program (Grant ID: 765148). S. L. was supported by the Faculty Research Grant (NO. FRG-18-020-FI) at Macau University of Science and Technology. B. W. and B. M. were supported through the DFG, SFB-824, subproject B12. K. C. was supported by Clinical Research Priority Program (CRPP) Grant on Artificial Intelligence in Oncological Imaging Network, University of Zurich.

## References

1. Aerts, H.J., et al.: Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nat. Commun.* **5**(1), 1–9 (2014)
2. Azizi, S., et al.: Big self-supervised models advance medical image classification. arXiv preprint [arXiv:2101.05224](https://arxiv.org/abs/2101.05224) (2021)
3. Bachman, P., Hjelm, R.D., Buchwalter, W.: Learning representations by maximizing mutual information across views. In: *Advances in Neural Information Processing Systems*, pp. 15535–15545 (2019)
4. Bakas, S., et al.: Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge. arXiv preprint [arXiv:1811.02629](https://arxiv.org/abs/1811.02629) (2018)
5. Chaitanya, K., Erdil, E., Karani, N., Konukoglu, E.: Contrastive learning of global and local features for medical image segmentation with limited annotations. arXiv preprint [arXiv:2006.10511](https://arxiv.org/abs/2006.10511) (2020)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. arXiv preprint [arXiv:2002.05709](https://arxiv.org/abs/2002.05709) (2020)
7. 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)
8. Chen, X., He, K.: Exploring simple siamese representation learning. arXiv preprint [arXiv:2011.10566](https://arxiv.org/abs/2011.10566) (2020)
9. Clark, K., et al.: The cancer imaging archive (TCIA): maintaining and operating a public information repository. *J. Digit. Imaging* **26**(6), 1045–1057 (2013). <https://doi.org/10.1007/s10278-013-9622-7>
10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: *Computer Vision and Pattern Recognition*, pp. 248–255 (2009)
11. Esteva, A., et al.: Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**(7639), 115–118 (2017)
12. Gillies, R.J., Kinahan, P.E., Hricak, H.: Radiomics: images are more than pictures, they are data. *Radiology* **278**(2), 563–577 (2016)
13. Grill, J.B., et al.: Bootstrap your own latent—a new approach to self-supervised learning. In: *Advances in Neural Information Processing Systems*, vol. 33 (2020)
14. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: *Computer Vision and Pattern Recognition*, pp. 1735–1742 (2006)
15. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Computer Vision and Pattern Recognition*, pp. 9729–9738 (2020)
16. Hofmanninger, J., Prayer, F., Pan, J., Röhrich, S., Prosch, H., Langs, G.: Automatic lung segmentation in routine imaging is primarily a data diversity problem, not a methodology problem. *Eur. Radiol. Exp.* **4**(1), 1–13 (2020). <https://doi.org/10.1186/s41747-020-00173-2>
17. Kiyasseh, D., Zhu, T., Clifton, D.A.: CLOCS: contrastive learning of cardiac signals. arXiv preprint [arXiv:2005.13249](https://arxiv.org/abs/2005.13249) (2020)
18. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
19. Menze, B.H., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans. Med. Imaging* **34**(10), 1993–2024 (2014)
20. van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint [arXiv:1807.03748](https://arxiv.org/abs/1807.03748) (2018)

21. Vu, Y.N.T., Wang, R., Balachandar, N., Liu, C., Ng, A.Y., Rajpurkar, P.: MedAug: contrastive learning leveraging patient metadata improves representations for chest X-ray interpretation. arXiv preprint [arXiv:2102.10663](https://arxiv.org/abs/2102.10663) (2021)
22. Ye, M., Zhang, X., Yuen, P.C., Chang, S.F.: Unsupervised embedding learning via invariant and spreading instance feature. In: Computer Vision and Pattern Recognition, pp. 6210–6219 (2019)
23. Yip, S.S., Aerts, H.J.: Applications and limitations of radiomics. *Phys. Med. Biol.* **61**(13), R150 (2016)
24. Zhou, Z., et al.: Models genesis: generic autodidactic models for 3D medical image analysis. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11767, pp. 384–393. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-32251-9\\_42](https://doi.org/10.1007/978-3-030-32251-9_42)
25. 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](https://doi.org/10.1007/978-3-030-32251-9_46)