

Integration of Self-Supervised BYOL in Semi-Supervised Medical Image Recognition

Hao Feng², Yuanzhe Jia², Ruijia Xu³, Mukesh Prasad¹, Ali Anaissi², and Ali Braytee^{1*}

¹ University of Technology Sydney, School of Computer Science, Ultimo, Australia

² The University of Sydney, School of Computer Science, Camperdown, Australia

³ North University of China, China

ali.braytee@uts.edu.au

Abstract. Image recognition techniques heavily rely on abundant labeled data, particularly in medical contexts. Addressing the challenges associated with obtaining labeled data has led to the prominence of self-supervised learning and semi-supervised learning, especially in scenarios with limited annotated data. In this paper, we proposed an innovative approach by integrating self-supervised learning into semi-supervised models to enhance medical image recognition. Our methodology commences with pre-training on unlabeled data utilizing the BYOL method. Subsequently, we merge pseudo-labeled and labeled datasets to construct a neural network classifier, refining it through iterative fine-tuning. Experimental results on three different datasets demonstrate that our approach optimally leverages unlabeled data, outperforming existing methods in terms of accuracy for medical image recognition.

Keywords: self-supervised learning · semi-supervised learning · medical image recognition · limited labels.

1 Introduction

Artificial Intelligence (AI) has significant potential in medical applications but still faces challenges in medical image recognition [1]. Medical image datasets are rich in spatial resolution, color channels, and diverse representations of organs, diseases, and anatomical structures. Obtaining relevant labels for medical classification is a primary hurdle, often requiring the expertise of medical professionals, leading to a time-consuming and logistically complex annotation process [2]. Semi-supervised learning offers a potential solution for limited labeled data. This approach can improve model performance by incorporating additional unlabeled data. Similarly, self-supervised learning has emerged as a robust strategy for exploiting the inherent patterns in unlabeled data and acquiring representations that capture the underlying semantics [3].

We propose a method to address the challenge of limited labeled data in medical image classification by combining the advantages of self-supervised and semi-supervised learning. In our method, Bootstrap Your Own Latent (BYOL) [4],

a pre-training model for semi-supervised learning is employed to acquire useful representations from large amounts of unlabeled data. These representations are then fine-tuned using smaller labeled datasets to achieve high performance. Our method offers two key advantages: firstly, BYOL enhances the model generalization by capturing structural and semantic information from unlabeled medical data. Secondly, semi-supervised learning optimizes the model performance by leveraging labeled and unlabeled data.

2 Related work

Medical image recognition heavily relies on labeled datasets. Many AI techniques, such as deep convolutional neural networks, perform optimally when abundant data and high-quality annotations are available [5]. Challenges persist when faced with limited labeled data [6]. To tackle this issue, unsupervised learning techniques like Sparse Coding (SC), Auto-encoders (AE), and Restricted Boltzmann Machines (RBMs) have been proposed. However, they are often limited to learning basic image features and struggle with the intricate characteristics present in medical images. Fortunately, leveraging unlabeled data through semi-supervised and self-supervised learning provides a promising avenue to reduce dependence on labeled data in medical image recognition.

Semi-supervised learning combines supervised and unsupervised methods, utilizing labeled and unlabeled data for enhanced performance. When dealing with brain images that lack precise diagnostic information, semi-supervised support vector machines provide feature labels to overcome classification challenges [7]. This approach is also effective on imbalanced data. Multi-Curriculum Pseudo-Labeling (MCPL) evaluates the learning progress and adjusts the threshold for each category. This adaptation improves the efficiency of semi-supervised learning, especially for imbalanced medical image classification [8]. Furthermore, semi-supervised learning can extract additional semantic information from unlabeled data [9]. Overall, by leveraging unlabeled data, semi-supervised learning yields valuable information and training examples, thereby reducing labeling costs, addressing data imbalances, and enhancing model performance.

Although semi-supervised learning has reduced the expenses associated with medical image labeling, it still necessitates labeled data, a formidable challenge in medical images. There is a pressing need to diminish the model's reliance on annotated data. Recent studies have shown that incorporating self-supervised learning into the anomaly detection framework can significantly improve the accuracy by exploiting valuable information from the original unlabeled data [10]. Moreover, a novel self-supervised learning algorithm, Bootstrap Your Own Latent (BYOL), was proposed and achieved a high classification accuracy on ImageNet, with the potential for further improvement using deeper architectures [4]. Unlike other self-supervised learning algorithms, such as Contrastive Predictive Coding (CPC) and SimCLR, BYOL distinguishes itself by eliminating the need to generate negative samples for comparison. Therefore, the BYOL algorithm is a promising choice for augmenting semi-supervised models.

3 A brief on BYOL

BYOL is a self-supervised learning method for feature extraction in semi-supervised learning tasks. One of its key merits lies in its independence from labeled data, as it exclusively leverages unlabeled data during the training process. This approach enables the model to derive meaningful representations from a large volume of unlabeled data, thereby enhancing its efficacy in semi-supervised learning applications.




Fig. 1. BYOL architecture overview [4]

As shown in Fig. 1, the core idea of BYOL is to develop two neural networks that learn from each other, the online network θ and the target network ξ , where v is the input of θ and v' is the input of ξ . v and v' are obtained by applying stochastic image augmentation methods t and t' on the input image x . f_θ and f_ξ are encoders that share the same structure with different parameters so that v and v' get their respective corresponding representations y_θ and y'_ξ after passing through f_θ and f_ξ , where y_θ is considered as a feature of the image that will be used eventually. After this, the online network θ and the target network ξ map y_θ and y'_ξ to the latent space to get z_θ and z'_ξ using g_θ and g'_ξ , respectively. The prediction result $q_\theta(z_\theta)$ can be obtained from the online network θ by adding a layer of network structure q_θ . The target of BYOL is trying to make $q_\theta(z_\theta)$ as close to z'_ξ as possible using the loss function in Eq. 1

$$\mathcal{L}_{\theta,\xi} \triangleq \left\| \bar{q}_\theta(z_\theta) - \bar{z}'_\xi \right\|_2^2 = 2 - 2 \cdot \frac{\langle q_\theta(z_\theta), z'_\xi \rangle}{\|q_\theta(z_\theta)\|_2 \cdot \|z'_\xi\|_2} \quad (1)$$

4 Proposed method

Our proposed method for medical image recognition is designed by utilising BYOL in a semi-supervised setting, which mainly consists of two steps: pre-training and fine-tuning.

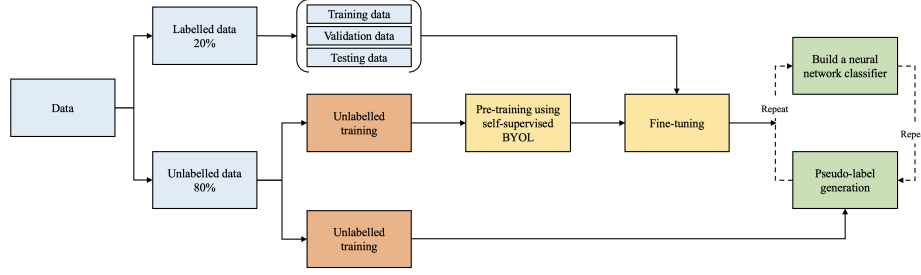


Fig. 2. The proposed semi-supervised learning with BYOL.

4.1 Pre-training

In our method, BYOL is employed to substitute the traditional pre-training process of semi-supervised learning, which relies on labeled data to generate the initial model. The optimisation process in Eq. 2, where η is the learning rate. The loss $L(\theta, \xi)$ is calculated and optimized at each training step, while the network parameters θ are dynamically updated by Stochastic Gradient Descent (SGD). The implementation of BYOL is relatively simple, with lower computing and memory requirements. As a result, we can streamline the initial phases of raw data preprocessing and feature extraction.

$$\theta \leftarrow optimizer(\theta, \nabla_{\theta} \mathcal{L}_{\theta, \xi}^{BYOL}, \eta), \quad \xi \leftarrow \tau \xi + (1 - \tau) \theta, \quad (2)$$

4.2 Fine-tuning

The feature extraction model generated by BYOL is fine-tuned with labeled data to construct a neural network classifier. This implementation is divided into three parts: constructing a neural network classifier, generating pseudo-labels, and model training.

4.2.1 Constructing a neural network classifier The pre-trained BYOL model is fine-tuned using a smaller dataset chosen for image classification. This entails substituting the prediction layer with one customized for the target task and applying supervised learning on labeled data. The network parameters are initialized with the pre-trained weights from the prior task, capitalizing on weight inheritance to enhance the initial state, accelerate training convergence, and exploit features acquired by the pre-trained model. The prediction layer is crafted with input dimensions aligned with the output of the projection layer in the online network and is concurrently updated with the online network. In addition, the loss function of BYOL transitions from image similarity to cross-entropy to provide intuitive insights and explanations. The cross-entropy loss function is defined in Eq. 3, where y represents the one-hot encoded ground truth labels for N samples, and p denotes the predicted class probabilities.

$$H(y, p) = - \sum_{i=1}^N y_i \log \left(\frac{e^{p_i}}{\sum_{j=1}^N e^{p_j}} \right) \quad (3)$$

4.2.2 Generating pseudo-labels Pseudo-labels are formulated by leveraging unlabeled data and generating predictions using the pre-trained model. These synthetic labels are then integrated into the original dataset, effectively updating the sample labels within the unlabeled data. Subsequently, the predicted labels are appended to the pseudo-label repository of the neural network classifier. This combination of pseudo-labels with a limited set of labeled data enlarges the dataset, encompassing both authentic and synthetic labels, albeit with potential noise. This expanded training dataset introduces a broader spectrum of data distributions, increased sample diversity, and enhanced informational content. Consequently, it empowers the model to acquire a more extensive understanding, elevating its capacity for generalization.

4.2.3 Model training The modified BYOL model undergoes retraining, leveraging both labeled and pseudo-labeled data. This fusion facilitates improved generalization and heightened accuracy in addressing the target task. Each iteration involves fine-tuning the model with labeled data, generating pseudo-labeled data based on the current model, and retraining the model with a combined dataset comprising both labeled and pseudo-labeled instances. This iterative optimization consistently elevates the model performance, enhancing its resilience and generalization capabilities. Striking a balance between the objectives of model improvement is also crucial throughout this process. Determining the optimal point for iterations, guided by specific requirements and constraints, is essential in selecting the most effective model.

5 Experiments and results

5.1 Datasets

Three datasets are used in our experiments: OCT2017 [11]; COVID-19 X-ray [12]; and Kvasir [13]. OCT2017⁴ consists of 109,312 images, of which 108,312 are used for training, 32 for validation and 968 for testing. COVID-19 X-ray⁵ contains 2,905 posterior-to-anterior chest X-ray images, including positive cases, viral pneumonia cases and normal cases. Kvasir⁶ collects manually annotated images covering several classes, including anatomical landmarks, pathological findings and endoscopic procedures, each class consisting of hundreds of images.

⁴ <https://tinyurl.com/OCT2017dataset>

⁵ <https://www.kaggle.com/datasets/tawsifurrahman/covid19-radiography-database>

⁶ <https://datasets.simula.no/kvasir/>

5.2 Hyperparameter tuning

Due to space constraints, we only demonstrate the tuning process on the OCT2017 dataset. Three hyperparameters are tested: the number of epochs, the learning rate, and the number of pseudo-labels. We first searched for the best combination of epochs and learning rates, exploring a range of epochs from 50 to 250 and a range of learning rates from 0.001 to 0.03. We employ accuracy as an evaluation metric. The model achieved its highest accuracy of 0.95 with 250 epochs and a learning rate of 0.001, although the computation time increased with a larger epoch count (see Table 1). We then searched for the optimal number of pseudo-labels using grid search with values of 500, 1000, and 2,000, which attained accuracy scores of 0.961, 0.962, and 0.966, respectively. This allowed us to establish the optimal hyperparameters for our model: 250 epochs, a learning rate of 0.001, and 2,000 pseudo-labels.

Table 1. Selecting the optimal epoch and learning rate on OCT2017.

Learning Rate/Epoch	50	150	250
0.03	0.51	0.86	0.83
0.01	0.62	0.81	0.90
0.001	0.83	0.93	0.95
Computation Time	15 min	42 min	94 min

5.3 Results

Table 2. Classification results using accuracy on three datasets.

Method/Dataset	OCT2017	COVID-19	Kvasir
MeanTeacher [14]	0.92	0.93	0.92
MixMatch [15]	0.93	0.94	0.91
FixMatch [16]	0.92	0.93	0.92
VAT [17]	0.92	0.91	0.91
VATNM [18]	0.93	0.92	0.92
GLM [19]	0.94	0.93	0.93
VTS [20]	0.95	0.95	0.93
NNM [21]	0.95	0.96	0.93
Ours	0.966	0.987	0.976

The experiments aim to evaluate the model performance across various datasets, including OCT2017, COVID-19 X-ray and Kvasir. Table 2 displays a comparative analysis of results across three distinct datasets, assessing the proposed model against various semi-supervised counterparts such as MeanTeacher, MixMatch, FixMatch, VAT, VATNM, GLM, VTS and NNM. Evaluation metrics

focus on test set accuracy. For OCT2017, benchmark models achieved scores between 0.92 and 0.95. Similarly, accuracy scores fell within the 0.91 to 0.96 range on COVID-19 X-rays, and performance was maintained within the 0.91 to 0.93 range on Kvasir. Noteworthy is the exceptional performance of our proposed method across all three datasets, achieving scores of 0.966 for OCT2017, 0.987 for COVID-19, and 0.976 for Kvasir. These results underscore the effectiveness and competitiveness of our approach when compared to existing methods.

6 Conclusion

This paper introduces a method to improve the performance of semi-supervised learning on medical image data by integrating self-supervised BYOL. Experimental findings demonstrate that BYOL markedly boosts semi-supervised learning, diminishing the dependence on labeled data and setting it apart from existing semi-supervised approaches.

Acknowledgement

We acknowledge Kexin Liu, Qian Deng, Ruyi Jin, Shuting Li, and Yongxin Dai for their invaluable contributions to the project.

References

1. Rajkomar, A., Dean, J., Kohane, I.: Machine learning in medicine. *New England Journal of Medicine* **380**, 1347–1358 (04 2019). <https://doi.org/10.1056/NEJMra1814259>
2. Litjens, G., Kooi, T., Bejnordi, B.E., Setio, A.A.A., Ciompi, F., Ghafoorian, M., van der Laak, J.A.W.M., van Ginneken, B., Sánchez, C.I.: A survey on deep learning in medical image analysis. *Medical Image Analysis* **42**, 60–88 (2017). <https://doi.org/10.1016/j.media.2017.07.005>
3. Chaturvedi, K., Braytee, A., Li, J., Prasad, M.: Ss-cpgan: Self-supervised cut-and-pasting generative adversarial network for object segmentation. *Sensors* **23**(7), 3649 (2023)
4. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Valko, M.: Bootstrap your own latent: A new approach to self-supervised learning. *arXiv.org* (2020)
5. Zhou, Y., Liu, J., Yang, Z., Liu, T., Meng, X., Zhou, Z., Anaissi, A., Braytee, A.: Vgg-fusionnet: a feature fusion framework from ct scan and chest x-ray images based deep learning for covid-19 detection. In: *2022 IEEE International Conference on Data Mining Workshops (ICDMW)*. pp. 1–9. IEEE (2022)
6. Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S.: Dermatologist-level classification of skin cancer with deep neural networks. *Nature (London)* **542**(7639), 115–118 (2017). <https://doi.org/10.1038/nature21056>
7. Filipovych, R., Davatzikos, C.: Semi-supervised pattern classification of medical images: Application to mild cognitive impairment (mci). *NeuroImage (Orlando, Fla.)* **55**(3), 1109–1119 (2011). <https://doi.org/10.1016/j.neuroimage.2010.12.066>

8. Peng, Z., Zhang, D., Tian, S., Wu, W., Yu, L., Zhou, S.K., Huang, S.: Faxmatch: Multi-curriculum pseudo-labeling for semi-supervised medical image classification. *Medical Physics (Lancaster)* **50**(5), 3210–3222 (2023). <https://doi.org/10.1002/mp.16312>
9. Liu, Q., Yu, L., Luo, L., Dou, Q., Heng, P.A.: Semi-supervised medical image classification with relation-driven self-ensembling model. *IEEE Transactions on Medical Imaging* **39**(11), 3429–3440 (2020). <https://doi.org/10.1109/TMI.2020.2995518>
10. Zhao, H., Li, Y., He, N., Ma, K., Fang, L., Li, H., Zheng, Y.: Anomaly detection for medical images using self-supervised and translation-consistent features. *IEEE Transactions on Medical Imaging* **40**(4), 1404–1415 (2021). <https://doi.org/10.1109/TMI.2020.3000458>
11. Kermany, D., Zhang, K., Goldbaum, M.: Labeled optical coherence tomography (oct) and chest x-ray images for classification. *Mendeley Data* (2018). <https://doi.org/10.17632/rscbjbr9sj.2>
12. Sahoo, P., Roy, I., Ahlawat, R., et al.: Potential diagnosis of covid-19 from chest x-ray and ct findings using semi-supervised learning. *Physics in Engineering and Science in Medicine* **45**(1), 31–42 (2022). <https://doi.org/10.1007/s13246-021-01075-2>
13. Liu, P., Qian, W., Cao, J., Xu, D.: Semi-supervised medical image classification via increasing prediction diversity. *Applied Intelligence (Dordrecht, Netherlands)* **53**(9), 10162–10175 (2023). <https://doi.org/10.1007/s10489-022-04012-2>
14. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems* **30** (2017)
15. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. *Advances in neural information processing systems* **32** (2019)
16. Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems* **33**, 596–608 (2020)
17. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* **41**(8), 1979–1993 (2018)
18. Cui, S., Wang, S., Zhuo, J., Li, L., Huang, Q., Tian, Q.: Towards discriminability and diversity: Batch nuclear-norm maximization under label insufficient situations. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 3941–3950 (2020)
19. Gyawali, P.K., Ghimire, S., Bajracharya, P., Li, Z., Wang, L.: Semi-supervised medical image classification with global latent mixing. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part I 23*. pp. 604–613. Springer (2020)
20. Wang, X., Chen, H., Xiang, H., Lin, H., Lin, X., Heng, P.A.: Deep virtual adversarial self-training with consistency regularization for semi-supervised medical image classification. *Medical image analysis* **70**, 102010 (2021)
21. Liu, P., Qian, W., Cao, J., Xu, D.: Semi-supervised medical image classification via increasing prediction diversity. *Applied Intelligence* **53**(9), 10162–10175 (2023)