



Contrastive Learning with Continuous Proxy Meta-Data for 3D MRI Classification

Benoit Dufumier, Pietro Gori, Julie Victor, Antoine Grigis, Michel Wessa, Paolo Brambilla, Pauline Favre, Mircea Polosan, Colm Mcdonald, Camille Marie Piguet, et al.

► To cite this version:

Benoit Dufumier, Pietro Gori, Julie Victor, Antoine Grigis, Michel Wessa, et al.. Contrastive Learning with Continuous Proxy Meta-Data for 3D MRI Classification. MICCAI, Sep 2021, Strasbourg (virtuel), France. hal-03262256

HAL Id: hal-03262256

<https://telecom-paris.hal.science/hal-03262256>

Submitted on 16 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Contrastive Learning with Continuous Proxy Meta-Data for 3D MRI Classification

Benoit Dufumier^{1,2}, Pietro Gori², Julie Victor¹, Antoine Grigis¹, Michel Wessa³, Paolo Brambilla⁴, Pauline Favre¹, Mircea Polosan⁵, Colm McDonald⁶, Camille Marie Piguet, and Edouard Duchesnay¹ for the Alzheimer’s Disease Neuroimaging Initiative

¹ NeuroSpin, CEA Saclay, Université Paris-Saclay, France
`benoit.dufumier@cea.fr`

² LTCI, Télécom Paris, IIParis, France

³ Dep. of Neuropsychology, Johannes-Gutenberg Univ. of Mainz, Germany

⁴ Dep. of Neurosciences, Fondazione IRCCS, University of Milan, Italy

⁵ Université Grenoble Alpes, Inserm U1216, CHU Grenoble Alpe, France

⁶ Centre for Neuroimaging & Cognitive Genomics (NICOG), Galway, Ireland

Abstract. Traditional supervised learning with deep neural networks requires a tremendous amount of labelled data to converge to a good solution. For 3D medical images, it is often impractical to build a large homogeneous annotated dataset for a specific pathology. Self-supervised methods offer a new way to learn a representation of the images in an unsupervised manner with a neural network. In particular, contrastive learning has shown great promises by (almost) matching the performance of fully-supervised CNN on vision tasks. Nonetheless, this method does not take advantage of available meta-data, such as participant’s age, viewed as prior knowledge. Here, we propose to leverage continuous *proxy* metadata, in the contrastive learning framework, by introducing a new loss called y -Aware InfoNCE loss. Specifically, we improve the positive sampling during pre-training by adding more positive examples with similar *proxy* meta-data with the anchor, assuming they share similar discriminative semantic features. With our method, a 3D CNN model pre-trained on 10^4 multi-site healthy brain MRI scans can extract relevant features for three classification tasks: schizophrenia, bipolar diagnosis and Alzheimer’s detection. When fine-tuned, it also outperforms 3D CNN trained from scratch on these tasks, as well as state-of-the-art self-supervised methods. Our code is made publicly available [here](#).

1 Introduction

Recently, self-supervised representation learning methods have shown great promises, surpassing traditional transfer learning from ImageNet to 3D medical images [29]. These models can be trained without costly annotations and they offer a great initialization point for a wide set of downstream tasks, avoiding the domain gap between natural and medical images. They mainly rely on a pretext task that is informative about the prior we have on the data. This proxy task essentially

consists in corrupting the data with non-linear transformations that preserve the semantic information about the images and learn the reverse mapping with a Convolutional Neural Network (CNN). Numerous tasks have been proposed both in the computer vision field (inpainting [21], localization of a patch [9], prediction of the angle of rotation [10], jigsaw [18], etc.) and also specifically designed for 3D medical images (context restoration [4], solving the rubik’s cube [30], sub-volumes deformation [29]). They have been successfully applied to 3D MR images for both segmentation and classification [24,26,29,30], outperforming the classical 2D approach with ImageNet pre-training. Concurrently, there has been a tremendous interest in contrastive learning [13] over the last year. Notably, this unsupervised approach almost matches the performance over fully-supervised vision tasks and it outperforms supervised pre-training [1,5,14]. A single encoder is trained to map semantically similar “positive” samples close together in the latent space while pushing away dissimilar “negative” examples. In practice, all samples in a batch are transformed twice through random transformations $t \sim \mathcal{T}$ from a set of parametric transformations \mathcal{T} . For a given reference point (anchor) x , the positive samples are the ones derived from x while the other samples are considered as negatives. Most of the recent works focus in finding the best transformations \mathcal{T} that degrade the initial image x while preserving the semantic information [5,27] and very recent studies intend to improve the negative sampling [6,22]. However, two different samples are not necessarily semantically different, as emphasized in [6,28], and they may even belong to the same semantic class. Additionally, two samples are not always equally semantically different from a given anchor and so they should not be equally distant in the latent space from this anchor. In this work, we assume to have access to continuous *proxy* meta-data containing relevant information about the images at hand (*e.g* the participant’s age). We want to leverage these *meta-data* during the contrastive learning process in order to build a more universal representation of our data. To do so, we propose a new *y*-Aware InfoNCE loss inspired from the Noise Contrastive Estimation loss [12] that aims at improving the positive sampling according to the similarity between two *proxy* meta-data. Differently from [17], i) we perform contrastive learning with continuous meta-data (not only categorical) and ii) our first purpose is to train a generic encoder that can be easily transferred to various 3D MRI target datasets for classification or regression problems in the very small data regime ($N \leq 10^3$). It is also one of the first studies to apply contrastive learning to 3D anatomical brain images [2]. Our main contributions are:

- we propose a novel formulation for contrastive learning that leverages *continuous* meta-data and derive a new loss, namely the *y*-Aware InfoNCE loss
- we empirically show that our unsupervised model pre-trained on a large-scale multi-site 3D brain MRI dataset comprising $N = 10^4$ healthy scans reaches or outperforms the performance of CNN model fully-supervised on 3 classification tasks under the linear protocol evaluation
- we also demonstrate that our approach gives better results when fine-tuning on 3 target tasks than training from scratch

- we finally performed an ablation study showing that leveraging the meta-data improves the performance for all the downstream tasks and different set of transformations \mathcal{T} compared to SimCLR [5]

2 Method

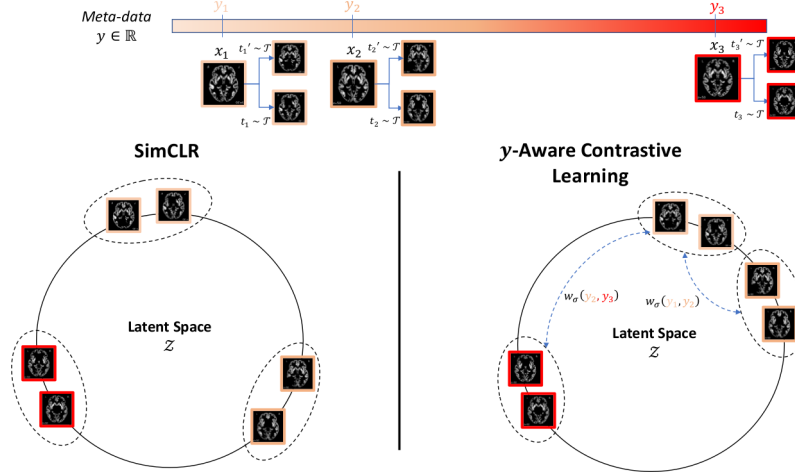


Fig. 1: Differently from SimCLR [5], our new loss can handle meta-data $y \in \mathbb{R}$ by redefining the notion of similarity between two images in the latent space \mathcal{Z} . For an image x_i , transformed twice through two augmentations $t_1, t_1' \sim \mathcal{T}$, the resulting views $(t_1(x_i), t_2(x_i))$ are expected to be close in the latent space through the learnt mapping f_θ , as in SimCLR. However, we also expect a different input $x_{k \neq i}$ to be close to x_i in \mathcal{Z} if the two proxy meta-data y_i and y_k are similar. We define a similarity function $w_\sigma(y_i, y_k)$ that quantifies this notion of similarity.

Problem Formalization. In contrastive learning [5, 17, 27], one wants to learn a parametric function $f_\theta : \mathcal{X} \mapsto \mathbb{S}^{d-1} = \mathcal{Z}$ between the input image space \mathcal{X} and the unit hypersphere, without meta-data. The goal of f_θ is to map samples to a representation space where semantically similar samples are “closer” than semantically different samples. To do that, each training sample $x_i \in \mathcal{X}$ is transformed twice through $t_1^i, t_2^i \sim \mathcal{T}$ to produce two augmented views of the same image $(v_1^i, v_2^i) := (t_1^i(x_i), t_2^i(x_i))$, where \mathcal{T} is a set of predefined transformations. Then, for each sample i , the model f_θ is trained to discriminate between the “positive” pair (v_1^i, v_2^i) , assumed to be drawn from the empirical joint distribution $p(v_1, v_2)$, and all the other “negative” pairs $(v_1^i, v_2^j)_{j \neq i}$, assumed to be drawn from the marginal distributions $p(v_1)p(v_2)$. With these assumptions, f_θ is an estimator of the mutual information I between v_1 and v_2 and it usually estimated by maximizing a lower bound of $I(v_1, v_2)$ called the InfoNCE loss [19]:

$$\mathcal{L}_{NCE} = -\log \frac{e^{f_\theta(v_1^i, v_2^i)}}{\frac{1}{n} \sum_{j=1}^n e^{f_\theta(v_1^i, v_2^j)}} \quad (1)$$

where n is the batch size, $f_\theta(v_1, v_2) := \frac{1}{\tau} f_\theta(v_1)^T f_\theta(v_2)$ and $\tau > 0$ is a hyperparameter. f_θ is usually defined as the composition of an encoder network $e_{\theta_1}(x)$ and a projection head z_{θ_2} (e.g. multi-layer perceptron) which is discarded after training (here $\theta = \{\theta_1, \theta_2\}$). Outputs lie on the unit hypersphere so that inner products can be used to measure cosine similarities in the representation space. In Eq. 1, every sample $v_2^j|_{j \neq i}$ is considered *equally* different from the anchor v_1^i . However, this is hardly true with medical images since we know, for instance, that two young healthy subjects should be considered more similar than a young and an old healthy subject. If we suppose to have access to continuous proxy metadata $y_i \in \mathbb{R}$ (e.g. participant's age or clinical score), then two views v_1^i, v_2^j with similar metadata y_i, y_j should be also close in the representation space \mathbb{S}^{d-1} . Inspired by vicinal risk minimization (VRM) [3], we propose to re-define $p(v_1, v_2)$ by integrating the proxy metadata y , modeled as a random variable, such that a small change in y results in a negligible change in $p(v_1, v_2|y)$. Similarly to [8], we define the empirical joint distribution as:

$$p_{emp}^{vic}(v_1, v_2|y) = \frac{1}{n} \sum_{i=1}^n \sum_{j=1}^n \frac{w_\sigma(y_i, y_j)}{\sum_{k=1}^n w_\sigma(y_i, y_k)} \delta(v_1 - v_1^i) \delta(v_2 - v_2^j) \quad (2)$$

where $\sigma > 0$ is the hyperparameter of the Radius Basis Function (RBF) kernel w_σ . Based on eq. 2, we can introduce our new y -Aware InfoNCE loss:

$$\mathcal{L}_{NCE}^y = -\sum_{k=1}^n \frac{w_\sigma(y_k, y_i)}{\sum_{j=1}^n w_\sigma(y_j, y_i)} \log \frac{e^{f_\theta(v_1^i, v_2^k)}}{\frac{1}{n} \sum_{j=1}^n e^{f_\theta(v_1^i, v_2^j)}} \quad (3)$$

In the limit case when $\sigma \rightarrow 0$, then we retrieve exactly the original InfoNCE loss, assuming that $y_i = y_j \Leftrightarrow x_i = x_j, \forall i, j \in [1..n]$. When $\sigma \rightarrow +\infty$, we assume that all samples $(x_i)_{i=1}^n$ belong to the same latent class.

Discrete case. If the proxy meta-data $(y_i)_{i \in [1..N]}$ are discrete, then we can simplify the above expression by imposing $w_\sigma(y_i, y_k) = \delta(y_i - y_k)$ retrieving the Supervised Contrastive Loss [17]. We may see \mathcal{L}_{NCE}^y as an extension of [17] in the continuous case. However, our purpose here is to build a robust encoder that can leverage meta-data to learn a more generalizable representation of the data.

Generalization. The proposed loss could be easily adapted to multiple metadata, both continuous and categorical, by defining one kernel per meta-data. Other choices of kernel, instead of the RBF, could also be considered.

Choice of the transformations \mathcal{T} . In our formulation, we did not specify particular transformations \mathcal{T} to generate (v_1^i, v_2^i) . While there have been recent works [4, 27] proposing transformations on natural images (color distorsion, cropping, cutout [7], etc.), there is currently no consensus for medical images in the context of contrastive learning. Here, we design three sets of transformations that preserve the semantic information in MR images: cutout, random cropping

and a combination of the two with also gaussian noise, gaussian blur and flip. Importantly, while color distortion is crucial on natural images [5] to avoid the model using a shortcut during training based on the color histogram, it is not necessarily the case for MR images (see Supp. 4).

3 Experiments

Datasets

- **Big Healthy Brains (BHB) dataset** We aggregated 13 publicly available datasets⁷ of 3D T1 MRI scans of healthy controls (HC) acquired on more than 70 different scanners and comprising $N = 10^4$ samples. We use this dataset only to pre-train our model with the **participant’s age as the *proxy* meta-data**. The learnt representation is then tested on the following four data-sets using as final task a binary classification between HC and patients.
- **SCHIZCONNECT-VIP**⁸ It comprises $N = 605$ multi-site MRI scans including 275 patients with strict schizophrenia (SCZ) and 330 HC.
- **BIOBD** [15, 23] This dataset includes $N = 662$ MRI scans acquired on 8 different sites with 356 HC and 306 patients with bipolar disorder (BD).
- **BSNIP** [25] It includes $N = 511$ MRI scans with $N = 200$ HC, $N = 194$ SCZ and $N = 117$ BD. This independent dataset is used only at test time in Fig. 2b).
- **Alzheimer’s Disease Neuroimaging Initiative (ADNI-GO)**⁹ We use $N = 387$ co-registered T1-weighted MRI images divided in $N = 199$ healthy controls and $N = 188$ Alzheimer’s patients (AD). We only included one scan per patient at the first session (baseline).

All data-sets have been pre-processed in the same way with a non-linear registration to the MNI template and a gray matter extraction step. The final spatial resolution is $1.5mm$ isotropic and the images are of size $121 \times 145 \times 121$.

Implementation details We implement our new loss based on the original InfoNCE loss [5] with Pytorch [20] and we use the Adam optimizer during training. As opposed to SimCLR [5] and in line with [2], we only use a batch size of $b = 64$ as it did not significantly change our results (see Supp. 4). We also follow [5] by fixing $\tau = 0.1$ in Eq.1 and Eq.3 and we set the learning rate to $\alpha = 10^{-4}$, decreasing it by 0.9 every 10 epochs. The model e_{θ_1} is based on a 3D adaptation of DenseNet121¹⁰ [16] and z_{θ_2} is a vanilla multilayer perceptron as in [5].

Evaluation of the representation In Fig.2, we compare the representation learnt using our model f_{θ} with the ones estimated using i) the InfoNCE loss (SimCLR) [5], ii) Model Genesis [29], a SOTA model for self-supervised learning with medical images, iii) a standard pre-training on age using a supervised approach (i.e. l_1 loss for age prediction), iv) BYOL [11] and MoCo [14] (memory

⁷ Demographic information as well as the public repositories can be found in Supp. 4

⁸ <http://schizconnect.org>

⁹ <http://adni.loni.usc.edu/about/adni-go>

¹⁰ Detailed implementation in our [repository](#)

bank $K = 1024$), 2 recently proposed SOTA models for representation learning, v) a multi-task approach SimCLR with age regression in the latent space (SimCLR+Age) and a fully fine-tuned supervised DenseNet trained to predict the final task. This can be considered as an upper bound, if the training dataset is sufficiently big. For the pre-training of our algorithm f_θ , we only use the BHB dataset with the participant’s age as *proxy* meta-data. For both contrastive learning methods and BYOL, we fix $\sigma = 5$ in Eq.3 and Eq.1 and only use random cutout for the transformations \mathcal{T} with a black patch covering $p = 25\%$ of the input image. We use UNet for pre-training with Model Genesis and DenseNet121 for all other models.

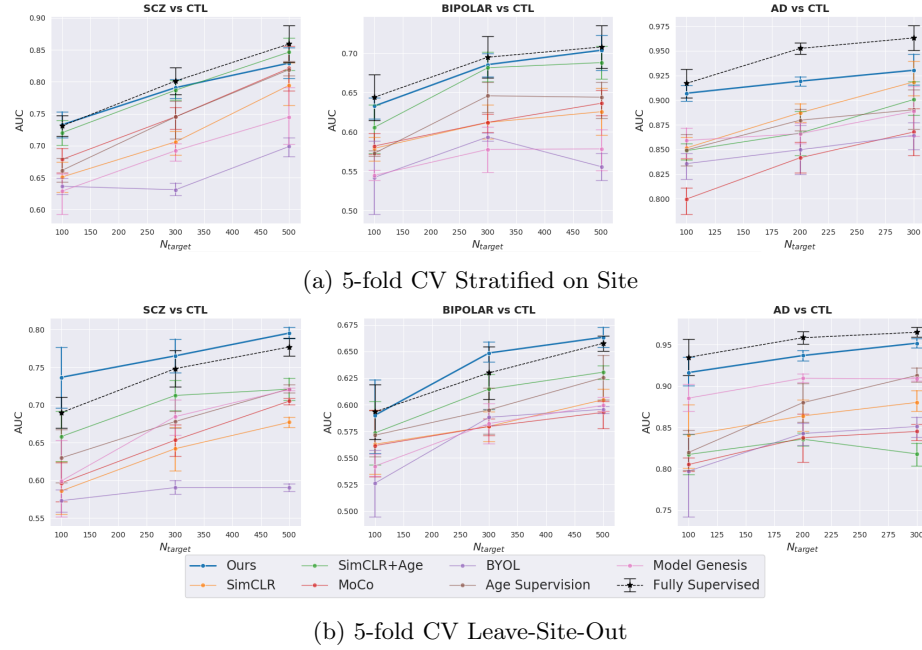


Fig. 2: Comparison of different representations in terms of classification accuracy (downstream task) on three different data-sets (one per column). Classification is performed using a linear layer on top of the pre-trained frozen encoders. (a) Data for training/validation and test come from the the same acquisition sites (b) Data for training/validation and test come from different sites.

In order to evaluate the quality of the learnt representations, we only added a linear layer on top of the frozen encoders pre-trained on BHB. We tune this linear layer on 3 different binary classification tasks (see Datasets section) with 5-fold cross-validation (CV). We tested two different situations: data for training/validation and test come either from the same sites (first row) or from different sites (second row). We also vary the size (i.e. number of subjects, N_{target})

of the training/validation set. For (a), we perform a stratified nested CV (two 5-fold CV, the inner one for choosing the best hyper-parameters and the outer one for estimating the test error). For (b), we use a 5-fold CV for estimating the best hyper-parameters and keep an independent constant test set for all N_{target} (see Supp. 4).

From Fig. 2, we notice that our method consistently outperforms the other pre-trainings even in the very small data regime ($N = 100$) and it matches the performance of the fully-supervised setting on 2 data-sets. Differently from age supervision, f_θ is less specialized on a particular proxy task and it can be directly transferred on the final task at hand without fine-tuning the whole network. Furthermore, compared to the multi-task approach SimCLR+Age, the features extracted by our method are less sensitive to the site where the MR images are coming from. This shows that our technique is the only one that efficiently uses the highly multi-centric dataset BHB by making the features learnt during pre-training less correlated to the acquisition sites.




Fig. 3: Linear classification performance on three binary classification tasks with $N_{pretrained} = 10^4$. *All TF* includes crop, cutout, gaussian noise, gaussian blur and flip. The encoder is frozen and we only tune a linear layer on top of it. $\sigma = 0$ corresponds to SimCLR [5] with InfoNCE loss. As we increase σ , we add more positive examples for a given anchor x_i with close proxy meta-data.

Importance of σ and \mathcal{T} in the positive sampling In Fig. 3, we study the impact of σ in Eq. 3 on the final representation learnt for a given set of transformations \mathcal{T} . As highlighted in [5], hard transformations seem to be important for contrastive learning (at least on natural images), therefore we have evaluated three different sets of transformations $\mathcal{T}_1 = \{ \text{Random Crop} \}$, $\mathcal{T}_2 = \{ \text{Random Cutout} \}$ and $\mathcal{T}_3 = \{ \text{Cutout, Crop, Gaussian Noise, Gaussian Blur, Flip} \}$. Importantly, we did not include color distortion in \mathcal{T}_3 since i) it is not adapted to MRI images where a voxel’s intensity encodes a gray matter density and ii) we did not observe significant difference between the color histograms of different scans as opposed to [5] (see Supp. 4). As before, we evaluated our representation

under the linear evaluation protocol. We can observe that \mathcal{T}_1 and \mathcal{T}_3 give similar performances with $\sigma > 0$, always outperforming both SimCLR ($\sigma = 0$) and age supervision on BHB. It also even outperforms the fully-supervised baseline on SCZ vs HC. We also find that a strong cropping or cutout strategy is detrimental for the final performances (see Supp. 4). Since \mathcal{T}_1 is computationally less expensive than \mathcal{T}_3 , we chose to use $\mathcal{T} = \mathcal{T}_1$ and $\sigma = 5$ in our experiments.

Fine-tuning Results Finally, we fine-tuned the whole encoder f_θ with different initializations on the 3 downstream tasks (see Table 1). To be comparable with Model Genesis [29], we also used the same UNet backbone for f_θ and we still fixed $\mathcal{T}_1 = \{\text{Random Cutout}\}$ and $\sigma = 5$. First, our approach outperforms the CNNs trained from scratch on all tasks as well as Model Genesis, even with the same backbone. Second, when using DenseNet, our pre-training remains better than using age supervision as pre-training for SCZ vs HC (even with the same transformations) and it is competitive on BD vs HC and AD vs HC.

Backbone	Pre-training	SCZ vs HC		BD vs HC		AD vs HC	
		$N_{train} = 100$	$N_{train} = 500$	$N_{train} = 100$	$N_{train} = 500$	$N_{train} = 100$	$N_{train} = 300$
UNet	None	72.62 \pm 0.9	76.45 \pm 2.2	63.03 \pm 2.7	69.20 \pm 3.7	88.12 \pm 3.2	94.16 \pm 3.9
	Model Genesis [29]	73.00 \pm 3.4	81.8 \pm 4.7	60.96 \pm 1.8	67.04 \pm 4.4	89.44 \pm 2.6	95.16 \pm 3.3
	SimCLR [4]	73.63 \pm 2.4	80.12 \pm 4.9	59.89 \pm 2.6	66.51 \pm 4.3	90.60 \pm 2.5	94.21 \pm 2.7
	Age Prediction w/ D.A	<u>75.32\pm2.2</u>	<u>85.27\pm2.3</u>	64.6\pm1.6	70.78\pm2.1	<u>91.71\pm1.1</u>	<u>95.26\pm1.5</u>
	Age-Aware Contrastive Learning (ours)	75.95\pm2.7	85.73\pm4.7	<u>63.79\pm3.0</u>	<u>70.35\pm2.7</u>	92.19\pm1.8	96.58\pm1.6
DenseNet	None	73.09 \pm 1.6	85.92 \pm 2.8	64.39 \pm 2.9	70.77 \pm 2.7	92.23 \pm 1.6	93.68 \pm 1.7
	None w/ D.A	<u>74.71\pm1.3</u>	86.94 \pm 2.8	64.79 \pm 1.3	72.25 \pm 1.5	92.10 \pm 1.8	94.16 \pm 2.5
	SimCLR [5]	70.80 \pm 1.9	86.35 \pm 2.2	60.57 \pm 1.9	67.99 \pm 3.3	91.54 \pm 1.9	94.26 \pm 2.9
	Age Prediction	72.90 \pm 4.6	<u>87.75\pm2.0</u>	64.60 \pm 3.6	72.07 \pm 3.0	92.07 \pm 2.7	<u>96.37\pm0.9</u>
	Age Prediction w/ D.A	74.06 \pm 3.4	86.90 \pm 1.6	65.79\pm2.0	<u>73.02\pm4.3</u>	94.01\pm1.4	96.10 \pm 3.0
	Age-Aware Contrastive Learning (ours)	76.33\pm2.3	88.11\pm1.5	<u>65.36\pm3.7</u>	73.33\pm4.3	<u>93.87\pm1.3</u>	96.84\pm2.3

Table 1: Fine-tuning results using 100 or 500 (300 for AD vs HC) training subjects. For each task, we report the AUC (%) of the fine-tuned models initialized with different approaches with 5-fold cross-validation. For age prediction, we employ the same transformations as in contrastive learning for the Data Augmentation (D.A) strategy. Best results are in **bold** and second bests are underlined.

4 Conclusion

Our key contribution is the introduction of a new contrastive loss, which leverages continuous (and discrete) meta-data from medical images in a self-supervised setting. We showed that our model, pre-trained with a large heterogeneous brain MRI dataset ($N = 10^4$) of healthy subjects, outperforms the other SOTA methods on three binary classification tasks. In some cases, it even reaches the performance of a fully-supervised network without fine-tuning. This demonstrates that our model can learn a meaningful and relevant representation of healthy brains which can be used to discriminate patients in small data-sets. An ablation study showed that our method consistently improves upon SimCLR for three different sets of transformations. We also made a step towards a debiased algorithm

by demonstrating that our model is less sensitive to the site effect than other SOTA fully supervised algorithms trained from scratch. We think this is still an important issue leading to strong biases in machine learning algorithms and it currently leads to costly harmonization protocols between hospitals during acquisitions. Finally, as a step towards reproducible research, we made our code public and we will release the BHB dataset to the scientific community soon. Future work will consist in developing transformations more adapted to medical images in the contrastive learning framework and in integrating other available meta-data (*e.g* participant’s sex) and modalities (*e.g* genetics). Finally, we envision to adapt the current framework for longitudinal studies (such as ADNI).

Acknowledgments: This work was granted access to the HPC resources of IDRIS under the allocation 2020-AD011011854 made by GENCI.

References

1. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. In: NeurIPS (2020)
2. Chaitanya, K., Erdil, E., Karani, N., Konukoglu, E.: Contrastive learning of global and local features for medical image segmentation with limited annotations (2020)
3. Chapelle, O., Weston, J., Bottou, L., Vapnik, V.: Vicinal risk minimization. MIT; 1998 (2001)
4. Chen, L., Bentley, P., Mori, K., Misawa, K., Fujiwara, M., Rueckert, D.: Self-supervised learning for medical image analysis using image context restoration. *Medical image analysis* **58**, 101539 (2019)
5. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2020)
6. Chuang, C.Y., Robinson, J., Lin, Y.C., Torralba, A., Jegelka, S.: Debiased contrastive learning. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) NeurIPS (2020)
7. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552* (2017)
8. Ding, X., Wang, Y., Xu, Z., Welch, W.J., Wang, Z.J.: CcGAN: Continuous conditional generative adversarial networks for image generation. In: ICLR (2021)
9. Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: ICCV (2015)
10. Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: ICLR (2018)
11. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., Piot, B., kavukcuoglu, k., Munos, R., Valko, M.: Bootstrap your own latent - a new approach to self-supervised learning. In: NeurIPS (2020)
12. Gutmann, M., Hyvärinen, A.: Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In: AISTATS (2010)
13. Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: CVPR (2006)
14. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020)

15. Hozer, F., Sarrazin, S., Laidi, C., Favre, P., Pauling, M., Cannon, D., McDonald, C., Emsell, L., Mangin, J.F., Duchesnay, E., et al.: Lithium prevents grey matter atrophy in patients with bipolar disorder: an international multicenter study. *Psychological medicine* pp. 1–10 (2020)
16. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR (2017)
17. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. *NeurIPS* (2020)
18. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: ECCV (2016)
19. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018)
20. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. In: *NeurIPS* (2019)
21. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: CVPR (2016)
22. Robinson, J., Chuang, C.Y., Sra, S., Jegelka, S.: Contrastive learning with hard negative samples. In: *ICLR* (2021)
23. Sarrazin, S., Cachia, A., Hozer, F., McDonald, C., Emsell, L., Cannon, D.M., Wessa, M., Linke, J., Versace, A., Hamdani, N., et al.: Neurodevelopmental subtypes of bipolar disorder are related to cortical folding patterns: An international multicenter study. *Bipolar disorders* **20**(8), 721–732 (2018)
24. Taleb, A., Loetzsch, W., Danz, N., Severin, J., Gaertner, T., Bergner, B., Lippert, C.: 3d self-supervised methods for medical imaging. In: *NeurIPS* (2020)
25. Tamminga, C.A., Pearlson, G., Keshavan, M., Sweeney, J., Clementz, B., Thaker, G.: Bipolar and schizophrenia network for intermediate phenotypes: outcomes across the psychosis continuum. *Schizophrenia bulletin* **40**, S131–S137 (2014)
26. Tao, X., Li, Y., Zhou, W., Ma, K., Zheng, Y.: Revisiting rubik’s cube: self-supervised learning with volume-wise transformation for 3d medical image segmentation. In: *MICCAI* (2020)
27. Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., Isola, P.: What makes for good views for contrastive learning? In: *NeurIPS* (2020)
28. Wei, C., Wang, H., Shen, W., Yuille, A.: Co2: Consistent contrast for unsupervised visual representation learning. In: *ICLR* (2021)
29. Zhou, Z., Sodha, V., Pang, J., Gotway, M.B., Liang, J.: Models genesis. *Medical Image Analysis* **67**, 101840 (2021)
30. 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: *MICCAI* (2019)

1. Demographic Information for BHB and UMAP visualization

Source	# Subjects	N	Age	Sex (%F)	# Sites
HCP	1113	1113	29 ± 4	45	1
IXI	559	559	48 ± 16	55	3
CoRR	1371	2897	26 ± 16	50	19
NPC	65	65	26 ± 4	55	1
NAR	303	323	22 ± 5	58	1
RBP	40	40	23 ± 5	52	1
OASIS 3	597	1262	67 ± 9	62	3
GSP	1570	1639	21 ± 3	58	1
ICBM	622	977	30 ± 12	45	3
ABIDE I	567	567	17 ± 8	17	20
ABIDE II	559	580	15 ± 9	30	17
Localizer	82	82	25 ± 7	56	2
MPI-Leipzig	316	316	37 ± 19	40	2
Total	7764	10420	32 ± 19	50	74

Fig. 4: Demographic information




Fig. 5: UMAP Representation

Fig. 4: Datasets aggregated for the BHB dataset. N is the number of scans that passed the QC. Note that there might be several sessions (scans) per subject. Fig. 5: 2D UMAP of ADNI features encoded (left) with SimCLR pre-training; (right) with our method. MRI from healthy participants with approximately the same age are mapped to the same region with our model.

2. Hyperparameter selection: Batch Size and Patch Size for Cutout and Crop

Fig. 6 (resp. 7) we reported the performance of our model pre-trained with varying batch size (resp. black patch size for cutout and crop size) and $\sigma = 5$. We performed a 5-fold CV under the linear evaluation protocol (encoder f_θ frozen) and we set $N_{target} = 500$ for SCZ vs HC and BIP vs HC and $N_{target} = 300$ for AD vs HC. We find that a large batch size is not necessarily required when dealing with brain MRI, in line with [2]. We also fixed $p = 25\%$ and $p' = 75\%$ in our study.

Batch Size	Target Task		
	SCZ vs HC	BIP vs HC	AD vs HC
64	82.94 \pm 2.7	70.36 \pm 2.6	93.03 \pm 1.8
100	84.15 \pm 2.7	70.42 \pm 1.1	93.53 \pm 1.6

Fig. 6: AUC score (%) as we vary the batch size during pre-training.

Transformations	Target Task		
	SCZ vs HC	BIP vs HC	AD vs HC
Cutout $p = 25\%$ $p = 50\%$	82.94 \pm 2.7	70.36 \pm 2.6	93.03 \pm 1.8
	84.00 \pm 2.1	68.96 \pm 2.2	89.21 \pm 2.7
Crop $p' = 75\%$ $p' = 50\%$	84.73 \pm 0.7	69.77 \pm 4.3	94.88 \pm 2.7
	81.77 \pm 3.1	68.69 \pm 1.3	91.46 \pm 3.2

Fig. 7: AUC score (%). The black patch size p (for random cutout) and the crop size p' are set during pre-training.

3. Histogram of colors for MR images

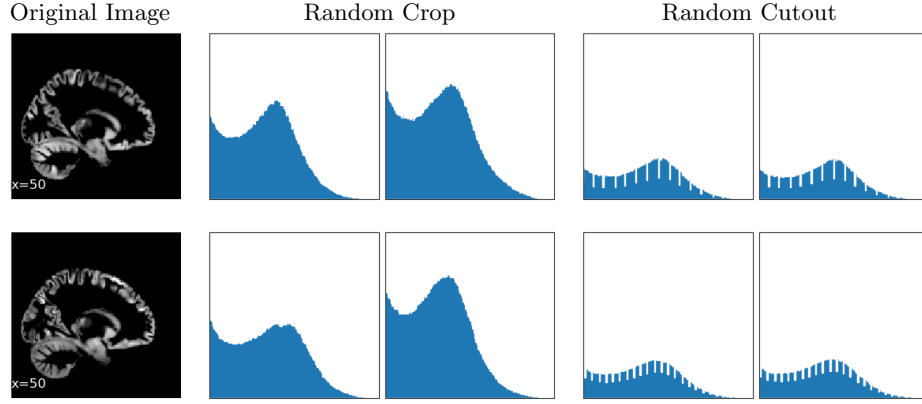


Fig. 8: Histogram of pixel intensities for 2 different images either i) randomly cropped or ii) partially masked with random cutout. We do not observe strong differences between the histograms for a given transformation. As such, color distortion may not be as critical as in [5] to learn a robust representation since the network cannot take a shortcut based only on the color histogram.

4. Cross-Validation Strategies




Fig. 9: Cross-Validation Strategies used when evaluating our representation on the 3 downstream tasks. We both tested (a) and (b) with varying N_{target} .