



**HAL**  
open science

# CONTRASTIVE LEARNING FOR REGRESSION IN MULTI-SITE BRAIN AGE PREDICTION

Carlo Alberto Barbano, Benoit Dufumier, Edouard Duchesnay, Marco  
Grangetto, Pietro Gori

► **To cite this version:**

Carlo Alberto Barbano, Benoit Dufumier, Edouard Duchesnay, Marco Grangetto, Pietro Gori. CONTRASTIVE LEARNING FOR REGRESSION IN MULTI-SITE BRAIN AGE PREDICTION. IEEE ISBI (International Symposium on Biomedical Imaging ), Apr 2023, Cartagena de Indias, Colombia. hal-03993158

**HAL Id: hal-03993158**

**<https://telecom-paris.hal.science/hal-03993158v1>**

Submitted on 16 Feb 2023

**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 FOR REGRESSION IN MULTI-SITE BRAIN AGE PREDICTION

Carlo Alberto Barbano<sup>\*†</sup> Benoit Dufumier<sup>\*‡</sup> Edouard Duchesnay<sup>‡</sup> Marco Grangetto<sup>†</sup> Pietro Gori<sup>\*</sup>

<sup>\*</sup> LTCI, Télécom Paris, IPParis <sup>†</sup> University of Turin <sup>‡</sup> NeuroSpin, CEA, Université Paris Saclay

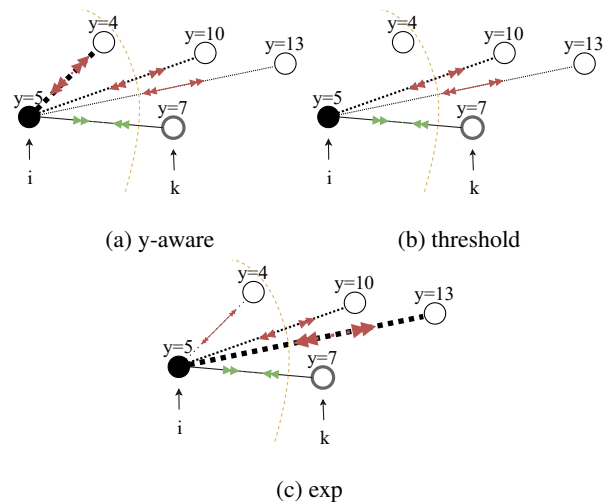
## ABSTRACT

Building accurate Deep Learning (DL) models for brain age prediction is a very relevant topic in neuroimaging, as it could help better understand neurodegenerative disorders and find new biomarkers. To estimate accurate and generalizable models, large datasets have been collected, which are often multi-site and multi-scanner. This large heterogeneity negatively affects the generalization performance of DL models since they are prone to overfit site-related noise. Recently, contrastive learning approaches have been shown to be more robust against noise in data or labels. For this reason, we propose a novel contrastive learning regression loss for robust brain age prediction using MRI scans. Our method achieves state-of-the-art performance on the OpenBHB challenge, yielding the best generalization capability and robustness to site-related noise.

**Index Terms**— MRI, multi-site, brain age, deep learning, contrastive learning, regression

## 1. INTRODUCTION

Brain aging involves complex biological processes, such as cortical thinning, that are highly heterogeneous across individuals, suggesting that people do not age in the same manner. Accurately modeling brain aging at the subject-level is a long-standing goal in neuroscience as it could enhance our understanding of age-related diseases such as neurodegenerative disorders. To this end, brain-age predictors linking neuroanatomy to chronological age have been proposed using Deep Learning (DL) [1]. In order to build accurate biomarker of aging, DL models need large-scale neuroimaging dataset for training, which often involves multi-site studies, partly because of the high cost per patient in each study. Recent works have shown that DL models, and in particular Deep Neural Networks (DNN), largely over-fit site-related noise when trained on such multi-site datasets, notably due to the difference in acquisition protocols, scanner constructors, physical properties such as permanent magnetic field [2, 3]. This also implies poor generalization performance on data from new incoming sites, highly limiting the applicability of these models to real-life scenarios. In order to build more robust and accurate brain age models insensitive to site, the OpenBHB chal-



**Fig. 1:** Comparison between different contrastive learning regression losses and their effect on the representations. Samples are aligned ( $\gg$   $\ll$ ) and repelled ( $\ll\gg$ ) with varying strength (line thickness) based on the continuous label  $y$ .

lenge [4] has been recently released. While most DNN used to derive brain age gap are usually trained as standard regressors with the optimization of mean absolute error [5, 6], Ridge or cross-entropy loss [1] (if age is binarized), these frameworks do not pay particular care about site-related information during training to produce robust representations of brain imaging data. On the other hand, contrastive learning paradigms for DNN training have been recently proposed in various contexts such as supervised [7], weakly-supervised [8, 9] and unsupervised representation learning [10, 11]. More importantly, contrastive learning has been shown to be more robust than traditional end-to-end approaches, such as cross-entropy, against noise in the data or the labels, resulting in better generalizing models [7, 12]. For this reason, in this work, we propose a novel contrastive learning loss for regression in the context of the OpenBHB challenge, where chronological age must be learned without being affected by site-related noise. With our method, we obtain the best results in the official leaderboard. Our contributions are:

- Novel contrastive learning regression loss for brain age prediction;
- State-of-the-art performance in brain age prediction on the

OpenBHB challenge.

## 2. METHOD

Supervised contrastive learning (i.e., SupCon [7]) leverages discrete labels (i.e., categories) to define positive and negative samples. Starting from a sample  $x_i$ , called the *anchor*, and its latent representation  $z_i = f(x_i)$ , SupCon aligns the representations of all positive samples (i.e. sharing the same class as  $x_i$ ) to  $z_i$ , while repelling the representations of the negative ones (i.e., different class). The SupCon loss is thus not adapted for regression (i.e., continuous labels), as it is not possible to determine a hard boundary between positive and negative samples. All samples are somehow positive and negative at the same time. Given the continuous label  $y_i$  for the anchor and  $y_k$  for a sample  $k$ , one could threshold the distance  $d$  between  $y_i$  and  $y_k$  at a certain value  $\tau$  in order to create positive and negative samples (i.e.,  $k$  is positive if  $d(y_i - y_k) < \tau$ ). The problem would then be how to choose  $\tau$ . Differently, we propose to define a degree of “positiveness” between samples using a kernel function  $w_k = K(y_i - y_k)$ , where  $0 \leq w_k \leq 1$ . Our goal is thus to learn a parametric function  $f : \mathcal{X} \rightarrow \mathbb{S}^d$  that maps samples with a high degree of positiveness ( $w_k \sim 1$ ) close in the latent space and samples with a low degree ( $w_k \sim 0$ ) far away from each other.

### 2.1. Contrastive Learning Regression Loss

In contrastive learning, we look for a mapping function  $f$  such that the following condition is always satisfied:  $s_t^- - s_k^+ \leq 0$ ,  $\forall t, k$ , where we denote as  $s_k = \text{sim}(f(x_i), f(x_k))$  the similarity between the representations of the  $k$ -th sample and the anchor  $i$  (e.g., cosine similarity) and  $s_t^-$  and  $s_k^+$  are the negative and positive samples respectively. The notion of “negative” (dissimilar from the anchor) and “positive” (similar to the anchor) samples is thus rooted in the contrastive learning framework. To adapt such framework for continuous labels, we propose to use a kernel function  $w_k$ , and we develop multiple formulations, illustrated in Fig. 1. To derive our proposed loss, we employ a metric learning approach, as in [11], which allows us to easily add conditioning and regularisation.

A first approach would be to consider as “positive” only the samples that have a degree of positiveness greater than 0, and align them with a strength proportional to the degree, namely:

$$\frac{w_k}{\sum_j w_j} (s_t - s_k) \leq 0 \quad \forall j, k, t \neq k \in A(i) \quad (1)$$

where we have normalized the kernel so that the sum over all samples is equal to 1 and we denote with  $A(i)$  the indices of samples in the minibatch distinct from  $x_i$ . Eq.1 can be transformed in an optimization problem using, as it is common in contrastive learning, the max operator and its smooth approx-

imation *LogSumExp*:

$$\begin{aligned} \arg \min_f \sum_k \max(0, \frac{w_k}{\sum_t w_t} \{s_t - s_k\}_{t=1, \dots, N, t \neq k}) = \\ \arg \min_f \sum_k \frac{w_k}{\sum_t w_t} \max(0, \{s_t - s_k\}_{t=1, \dots, N, t \neq k}) \quad (2) \\ \approx \mathcal{L}^{y\text{-aware}} = - \sum_k \frac{w_k}{\sum_t w_t} \log \left( \frac{\exp(s_k)}{\sum_{t=1}^N \exp(s_t)} \right) \end{aligned}$$

Interestingly, this is exactly the y-aware loss proposed in [13] for classification with weak continuous attributes. Due to the non-hard boundary between positive and negative samples, both  $s_t$  and  $s_k$  are defined over the entire minibatch. The kernel  $w_k$  is used to avoid aligning samples not similar to the anchor (i.e.  $w_k \approx 0$ ). It can be noticed that, while the nominator aligns  $x_k$ , in the denominator, the uniformity term (as defined in [14]) focuses more on the closest samples in the representation space: this could be undesirable, as these samples might have a greater degree of positiveness than the considered  $x_k$  (Fig. 1a).


To avoid that, we formulate a first extension ( $\mathcal{L}^{thr}$ ) of (1), which limits the uniformity term (i.e., denominator) to the samples that are more distant from the anchor than the considered  $x_k$  in the kernel space (omitting the normalization in the starting condition):

$$\begin{aligned} w_k (s_t - s_k) \leq 0 \quad \text{if } w_t - w_k \leq 0 \quad \forall k, t \neq k \in A(i) \\ \mathcal{L}^{thr} = - \sum_k \frac{w_k}{\sum_t \delta_{w_t < w_k} w_t} \log \left( \frac{\exp(s_k)}{\sum_{t \neq k} \delta_{w_t < w_k} \exp(s_t)} \right) \quad (3) \end{aligned}$$

Ideally,  $\mathcal{L}^{thr}$  avoids repelling samples more similar than  $x_k$ . However, it still focuses more on the closest sample “less positive” than  $x_k$ , i.e.  $x_t$  s.t.  $w_t > w_x$  and  $w_t \leq w_j \forall j \neq k$  (Fig. 1b). As noted in [11, 7], increasing the margin with respect to the closest “negative” sample works well for classification, however, it might not be best suited for regression. For this reason, we propose a second formulation ( $\mathcal{L}^{exp}$ ) that takes an opposite approach. Instead of focusing on repelling the closest “less positive” sample, we increase the repulsion strength for samples proportionally to their distance from the anchor in the kernel space:

$$\begin{aligned} w_k [s_t (1 - w_t) - s_k] \leq 0 \quad \forall k, t \neq k \in A(i) \\ \mathcal{L}^{exp} = - \frac{1}{\sum_t w_t} \sum_k w_k \log \frac{\exp(s_k)}{\sum_{t \neq k} \exp(s_t (1 - w_t))} \quad (4) \end{aligned}$$

In the resulting  $\mathcal{L}^{exp}$  formulation, the weighting factor  $1 - w_t$  acts like a temperature value, by giving more weight to the samples which are farther away from the anchor in the kernel space (Fig. 1c). Also, for a proper kernel choice, samples closer than  $x_k$  will be repelled with very low strength ( $\sim 0$ ). We argue that this approach is more suited for continuous attributes (i.e., regression task), as it enforces that samples close in the kernel space will be close in the representation space.



**Fig. 2:** Ablation study of the kernel functions. The Gaussian kernel (rbf) with  $\sigma = 2$  yields the best generalization results (Ext. MAE) and final score across all three loss functions. We also notice an overall slight improvement in the site balanced accuracy.

### 3. EXPERIMENTAL DATA

We conduct our experiments on the OpenBHB dataset, which contains 5330 3D brain MRI scans from 71 different acquisition sites. Every scan comes from a different subject. The subjects have European-American, European and Asian origins, to promote diversity. Three modalities are available, derived from the same T1-w MRI scan (VBM, SBM, and quasi-rar). We focus this study on gray matter volumes (VBM). The model evaluation on OpenBHB is done on two private test sets (internal and external). The internal test set contains the same sites as training, the external contains unseen ones.

### 4. EXPERIMENTS AND RESULTS

As network architecture, we employ the 3D implementations of ResNet-18 (33.2M parameters), AlexNet (2.5M parameters), and DenseNet-121 (11.3M parameters). For comparison with [4], we use the Adam optimizer, with an initial learning rate of  $10^{-4}$  decayed by a factor of 0.9 every 10 epochs, and a weight decay of  $5 * 10^{-5}$ . We use a batch size of 32, and train for a total of 300 epochs. Our trainings are implemented in PyTorch, and run on an NVIDIA A40 GPU, with a single training taking  $\sim 24$ h. For every model, we evaluate the mean absolute error (MAE) on both the internal and external test sets. Furthermore, the OpenBHB challenge computes a balanced accuracy (BAcc) for site prediction, training a logistic regression on the model representations. The final challenge score is then computed as  $\mathcal{L}_c = \text{BAcc}^{0.3} \cdot \text{MAE}_{\text{ext}}$ .

#### 4.1. Kernel function ablation study

We test two different kernels: a Gaussian kernel  $K_g(u) = \exp(-||u||^2/2\sigma^2)$  and Cauchy kernel  $K_c(u) = 1/(\gamma||u||^2 + 1)$ . We perform an ablation study for the two different kernels and hyperparameters, employing a ResNet-18 model. Fig. 2 shows the result. For each kernel, we report the metrics on the test set along with the final challenge score, for the three loss functions. Focusing on the final score, it's easy to see

Kernel	$\sigma / \gamma$	$\mathcal{L}^{y\text{-aware}}$	$\mathcal{L}^{\text{threshold}}$	$\mathcal{L}^{\text{exp}}$
Cauchy	1	2.15	2.28	1.82
	2	2.48	3.03	1.83
RBF	1	2.43	2.63	1.58
	2	<b>1.82</b>	<b>1.74</b>	<b>1.54</b>

**Table 1:** Ablation study of kernel functions, in terms of challenge's score.

that a Gaussian kernel with  $\sigma = 2$  produces the best results for all losses (for readability, the final score is also reported in Tab. 1). This can be attributed to the overall lower error on the external set (Ext. MAE), showing that, with this setting, the models can generalize better. Furthermore, we also notice an overall lower balanced accuracy for site prediction, showing that this configuration is somewhat more robust to site noise.

#### 4.2. Comparison of contrastive regression losses

In Tab. 2 we compare the results obtained with the different losses. Focusing on the aggregate score, the best results are obtained with  $\mathcal{L}^{\text{exp}}$  (1.54). Furthermore,  $\mathcal{L}^{\text{exp}}$  also outperforms the other losses in every evaluated metric. Most significantly, it shows the best generalization capability in the external test set, which, undoubtedly, is the most relevant result from a practical clinical perspective. On the internal test, we score a MAE of 2.55, which is also slightly better than the related literature on a similarly sized dataset with UKB [1]. Interestingly,  $\mathcal{L}^{\text{exp}}$  also shows the best robustness to site-related noise (with a BAcc of 5.1), which indicates that the learned space preserves the neuroanatomical features very well while also removing site noise.

#### 4.3. Final results on the OpenBHB challenge

Finally, we report the ranking of  $\mathcal{L}^{\text{exp}}$  of the OpenBHB leaderboard, testing also AlexNet and DenseNet-121. Tab. 3

Method	Int. MAE	BAcc	Ext. MAE	$\mathcal{L}_c$
$\mathcal{L}^{y\text{-aware}}$	2.66±0.00	6.60±0.17	4.10±0.01	1.82
$\mathcal{L}^{thr}$	2.95±0.01	5.73±0.15	4.10±0.01	1.74
$\mathcal{L}^{exp}$	<b>2.55±0.00</b>	<b>5.1±0.1</b>	<b>3.76±0.01</b>	<b>1.54</b>

**Table 2:** Comparison of contrastive losses.

Method	Model	Int. MAE	BAcc	Ext. MAE	$\mathcal{L}_c$
Baseline ( $\ell_1$ )	DenseNet	2.55±0.01	8.0±0.9	7.13±0.05	3.34
	ResNet-18	2.67±0.05	6.7±0.1	4.18±0.01	1.86
	AlexNet	2.72±0.01	8.3±0.2	4.66±0.05	2.21
ComBat	DenseNet	5.92±0.01	2.23±0.06	10.48±0.17	3.38
	ResNet-18	4.15±0.01	<b>4.5±0.0</b>	4.76±0.03	1.88
	AlexNet	3.37±0.01	6.8±0.3	5.23±0.12	2.33
$\mathcal{L}^{exp}$	DenseNet	2.85±0.00	5.34±0.06	4.43±0.00	1.84
	ResNet-18	<b>2.55±0.00</b>	5.1±0.1	<b>3.76±0.01</b>	<b>1.54</b>
	AlexNet	2.77±0.01	5.8±0.1	4.01±0.01	1.71

**Table 3:** Final scores on the OpenBHB leaderboard. Reference results from [4].

shows the results. We compare with baseline models [4] trained with the L1 loss, and with ComBat [15], a site harmonization algorithm developed for MRIs. Our proposed  $\mathcal{L}^{exp}$  achieves state-of-the-art performance on the final leaderboard, scoring the best final score and metrics on both the internal and external test set, with ResNet-18. The improvement observed in the external test is also reflected for both AlexNet and DenseNet compared to all baselines. For these models, the internal MAE reached by the L1 baseline is slightly lower than  $\mathcal{L}^{exp}$ . However, when looking at the other metrics, it is easy to see that this is due to more overfitting on the internal sites for the baseline. Lastly, regarding the balanced accuracy, we observe a significant improvement with respect to the L1 baseline, showing that  $\mathcal{L}^{exp}$  possesses some debiasing capability towards site noise. Besides AlexNet, however, ComBat still achieves a lower accuracy, showing that there is room for improvement.

## 5. CONCLUSIONS AND FUTURE WORKS

To the best of our knowledge, this is the first work proposing a contrastive learning loss for regression in medical imaging. We employ it to predict chronological brain age, using the OpenBHB multi-site MRI challenge. We achieve state-of-the-art performance on the challenge score and in terms of generalization capabilities to unseen sites. This represents a valuable result for the scientific community, as building a robust brain age prediction model can help in better understanding neurodevelopmental or neurodegenerative disorders. We empirically demonstrate that contrastive learning is also capable of (partially) debiasing the site effects. In future, we will extend our experiments to include debiasing techniques, such

as regularization terms [11] adapted for regression, to further improve the results.

**Compliance with ethical standards** This study was conducted using human subject data, available in open access. Ethical approval was not required.

## 6. REFERENCES

- [1] Han Peng et al., “Accurate brain age prediction with lightweight deep neural networks,” *MedIA*, 2021.
- [2] Christian Wachinger et al., “Detect and correct bias in multi-site neuroimaging datasets,” *MedIA*, 2021.
- [3] Ben Glocker et al., “Machine learning with multi-site imaging data: An empirical study on the impact of scanner effects,” in *MedNeurIPS Workshop*, 2019.
- [4] Benoit Dufumier et al., “Openbhb: a large-scale multi-site brain mri data-set for age prediction and debiasing,” *NeuroImage*, 2022.
- [5] James H Cole et al., “Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker,” *NeuroImage*, 2017.
- [6] B. A. Jonsson et al., “Brain age prediction using deep learning uncovers associated sequence variants,” *Nature Communications*, vol. 10, no. 1, Nov. 2019.
- [7] Prannay Khosla et al., “Supervised contrastive learning,” in *NeurIPS*, 2020.
- [8] Yao-Hung Hubert Tsai et al., “Conditional contrastive learning with kernel,” *arXiv:2202.05458*, 2022.
- [9] Benoit Dufumier et al., “Rethinking positive sampling for contrastive learning with kernel,” *arXiv:2206.01646*, 2022.
- [10] Ting Chen et al., “A simple framework for contrastive learning of visual representations,” in *ICML*, 2020.
- [11] Carlo Alberto Barbano et al., “Unbiased supervised contrastive learning,” *arXiv:2211.05568*, 2022.
- [12] Florian Graf et al., “Dissecting supervised contrastive learning,” in *ICML*, 2021.
- [13] Benoit Dufumier et al., “Contrastive learning with continuous proxy meta-data for 3d mri classification,” in *MICCAI*. Springer, 2021.
- [14] Tongzhou Wang and Phillip Isola, “Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere,” *ICML*, 2020.
- [15] Jean Philippe Fortin et al., “Harmonization of cortical thickness measurements across scanners and sites,” *NeuroImage*, 2018.