



**HAL**  
open science

## Optimizing contrastive learning for cortical folding pattern detection

Aymeric Gaudin, Louise Guillon, Clara Fischer, Arnaud Cachia, Denis Rivière, Jean-François Mangin, Joel Chavas

► **To cite this version:**

Aymeric Gaudin, Louise Guillon, Clara Fischer, Arnaud Cachia, Denis Rivière, et al.. Optimizing contrastive learning for cortical folding pattern detection. SPIE Medical Imaging 2024, Feb 2024, San Diego, United States. pp.12926-22. hal-04420177

**HAL Id: hal-04420177**

**<https://hal.science/hal-04420177>**

Submitted on 26 Jan 2024

**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.



Distributed under a Creative Commons Attribution 4.0 International License

# Optimizing contrastive learning for cortical folding pattern detection

Aymeric Gaudin<sup>a</sup>, Louise Guillon<sup>a</sup>, Clara Fischer<sup>a</sup>, Arnaud Cachia<sup>b</sup>, Denis Rivière<sup>a</sup>,  
Jean-François Mangin<sup>a</sup>, and Joël Chavas<sup>a</sup>

<sup>a</sup>Neurospin, Bâtiment 145, CEA Saclay, 91191 Gif-sur-Yvette, France

<sup>b</sup>LaPsyDé, Laboratoire A.Binet-Sorbonne, 46 rue Saint Jacques, 75005 Paris

## ABSTRACT

The human cerebral cortex has many bumps and grooves called gyri and sulci. Even though there is a high inter-individual consistency for the main cortical folds, this is not the case when we examine the exact shapes and details of the folding patterns. Because of this complexity, characterizing the cortical folding variability and relating them to subjects' behavioral characteristics or pathologies is still an open scientific problem. Classical approaches include labeling a few specific patterns, either manually or semi-automatically, based on geometric distances, but the recent availability of MRI image datasets of tens of thousands of subjects makes modern deep-learning techniques particularly attractive. Here, we build a self-supervised deep-learning model to detect folding patterns in the cingulate region. We train a contrastive self-supervised model (SimCLR) on both Human Connectome Project (1101 subjects) and UKBioBank (21070 subjects) datasets with topological-based augmentations on the cortical skeletons, which are topological objects that capture the shape of the folds. We explore several backbone architectures (convolutional network, DenseNet, and PointNet) for the SimCLR. For evaluation and testing, we perform a linear classification task on a database manually labeled for the presence of the "double-parallel" folding pattern in the cingulate region, which is related to schizophrenia characteristics. The best model, giving a test AUC of 0.76, is a convolutional network with 6 layers, a 10-dimensional latent space, a linear projection head, and using the branch-clipping augmentation. This is the first time that a self-supervised deep learning model has been applied to cortical skeletons on such a large dataset and quantitatively evaluated. We can now envisage the next step: applying it to other brain regions to detect other biomarkers. The GitHub repository is publicly available on [https://github.com/neurospin-projects/2022\\_jchavas.cingulate.inhibitory\\_control](https://github.com/neurospin-projects/2022_jchavas.cingulate.inhibitory_control).

**Keywords:** Contrastive learning, SimCLR, MRI, brain, folding pattern

## 1. INTRODUCTION

The human cerebral cortex is folded and made of gyri separated by folds. Even though the main folding patterns are common across humans, the details of the cortical folding present a very high inter-subject variability, just like fingerprints. This variability is both a challenge and an opportunity: studying it can let us find anatomical biomarkers of neurodevelopmental pathologies or of subjects' endophenotypes linked to genetic polymorphisms.

We can decipher this variability by looking for local shape features called cortical folding patterns or folding patterns.

The literature shows few cases of cortical folding patterns that have been identified and correlated with pathologies or functions. For example, in the precentral region, Mellerio et al. (2015) manually identified a pattern named the Power Button Sign (PBS) pattern because of its recognizable shape (Fig. 1A). This pattern is very rare in the general population but is present in about 60% of patients with dysplasia near the motor area, leading to epilepsy.<sup>1</sup> Also, the Anterior Cingulate Cortex (ACC) has two identified folding patterns, the single-fold pattern and the double-parallel pattern (Fig. 1B):<sup>3</sup> the presence of different ACC patterns in both hemispheres is related to control efficiency in preschoolers,<sup>2</sup> and the shape of the ACC folds has been related to hallucinations among human patients suffering from schizophrenia.<sup>4</sup> To complement the manual methods, Borne et al. proposed a supervised model based on a convolutional network to detect known patterns.<sup>5</sup>

These manual techniques have a drawback: they are long to implement, while the supervised ones will not permit the detection of new folding patterns. We propose to resort to unsupervised techniques that allow the use

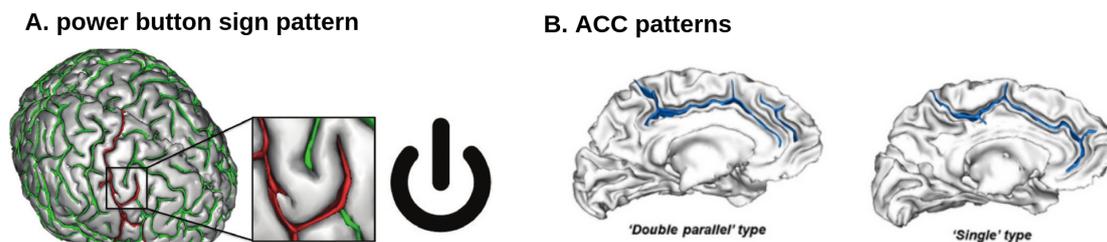


Figure 1. *Examples of cortical patterns.* **A.** The power button sign pattern in the central region is linked to epilepsy.<sup>1</sup> **B.** The anterior cingulate cortex (ACC) has two known patterns, a double-parallel pattern and a single-fold pattern.<sup>2</sup> We use the detection of these two ACC patterns to evaluate our models.

of the large datasets that are now available. Recently,  $\beta$ -VAE<sup>6</sup> and SimCLR<sup>7</sup> models have been trained to build a latent space for folding representation.<sup>8-10</sup> However, these studies have not used a manually labeled dataset for hyperparameter tuning and evaluation.

Here, we build a contrastive learning framework to find a representation that permits the automatic detection of relevant folding patterns. We use as inputs the so-called cortical skeletons—topologically-defined surfaces that follow the middle of a sulcus (Fig 3, blue surface in the left image)—, which are stable throughout life after early childhood and are expected to be less sensitive to data acquisition sites. We optimize a SimCLR model and evaluate it by detecting the double-parallel pattern in the anterior cingulate cortex of the right hemisphere.

## 2. METHODS

### 2.1 Datasets

Three datasets are used in this project: the human connectome project (HCP) dataset, the UKBioBank dataset, and the so-called ACC dataset :

- Data collection and sharing for the HCP dataset was provided by the MGH-USC Human Connectome Project.<sup>11</sup> We used in this work MRIs of 1101 HCP subjects,
- We preprocessed 21070 subjects of the UKBioBank, which recruited 500,000 people,<sup>12</sup>
- The ACC dataset stands for the Anterior Cingulate Cortex dataset. This dataset contains MRIs from 341 young subjects taken from 4 separate studies.<sup>13-17</sup> Each MRI has been manually labeled for the presence or absence of the paracingulate fold, which lies parallel above the anterior cingulate fold.<sup>3,18</sup> The presence of the paracingulate fold corresponds to what we call in this paper the double-parallel pattern.

The HCP dataset is randomly split in two, HCP-1 and HCP-2. The ACC dataset is also randomly split in two, ACC-1 for parameter tuning and ACC-2 for testing, using a stratification guaranteeing equipartition of the sites (as there can be an effect of the site on MRI data acquisition), gender, and subjects having a paracingulate fold in the right hemisphere (Fig. 2).

### 2.2 Demographic information

The 341 subjects of the ACC dataset are distributed as follows: 197 subjects are from a longitudinal study of childhood-onset schizophrenia,<sup>13,14</sup> out of which 73 are schizophrenia patients, 49 are patients' siblings, and 75 are healthy subjects.<sup>15</sup> 65 subjects (46 adolescents, 19 children) are taken from Delalande et al., 2020.<sup>16</sup> Last, 79 subjects are taken from Tissier et al., 2018.<sup>17</sup> All 341 subjects have been labeled manually for the presence of the paracingulate fold.

Gender distribution is given in Table. 1. HCP subjects were recruited from Missouri (United States), UK-BioBank subjects in 2006-2010 from across the UK, whereas ACC subjects were from France and the USA (Table 1).

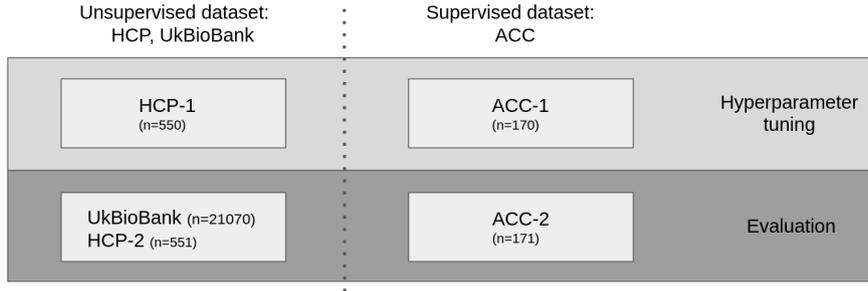


Figure 2. Dataset splits used for parameter tuning and evaluation.

	#subjects	Age Mean [min, max]	% of male	Origin	#sites
UKB	21070	64 [44, 82]	47	UK	1
HCP	1101	29 [22, 40]	48	USA	1
ACC	341	15 [8, 40]	58	France, USA	4

Table 1. Demographic information of used datasets.

### 2.3 Preprocessing

All subjects’ brain magnetic resonance images (MRIs), affinely normalized to a standard brain referential (ICBMc2009), were preprocessed with the Morphologist pipeline from the BrainVISA software\* to get their folding graph (Fig. 3), namely a graph-based representation of the cortical skeleton. Each branch of the graph is either a simple surface, a junction line between simple surfaces, or a bottom line of a fold; they are connected in the graph if their corresponding surfaces touch each other.<sup>19,20</sup>

\*<https://brainvisa.info>

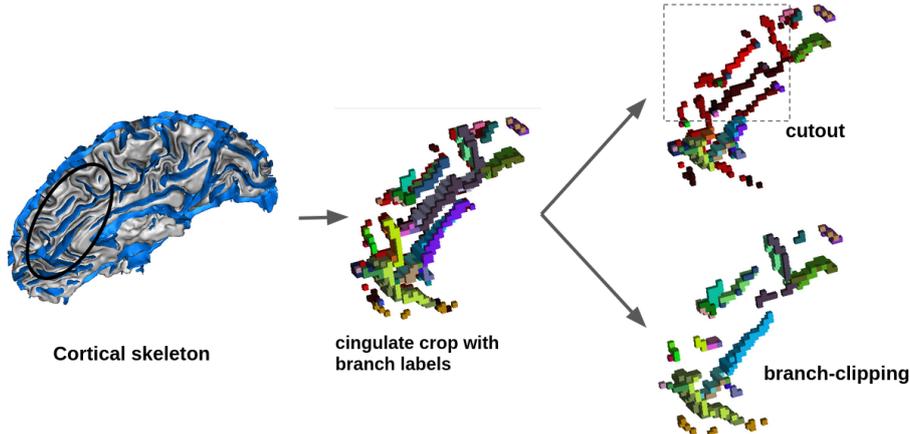


Figure 3. Preprocessing the inputs of the deep learning algorithm. **Left.** The cortical skeleton (in blue) represents the cortical folds for the right hemisphere. **Middle.** The cingulate crop with branch labels (one color per branch label): each branch represents a simple surface, a junction line, or a bottom line of a fold. **Right.** From the cingulate crop, we construct random views, using either the cutout augmentation, in which only bottom branches (represented in red) are kept inside the cutout (top right), or the branch-clipping augmentation, in which branches are randomly removed, and all bottom branches are removed (bottom right) .

We focus our study on the cingulate region of the right hemisphere (Fig. 3). For this, we compute a mask of the cingulate region over a database where the folds were manually labeled.<sup>21</sup> Our final inputs are 2-mm resolution 3D crops of dimension 17x40x38 (Fig. 3, middle).

## 2.4 Model

The contrastive model used is SimCLR, an instance discrimination contrastive model.<sup>7</sup> For each input image  $x$  of a batch of size  $N$ , we generate two views,  $x_i$  and  $x_j$ , at each epoch, whose model outputs are  $z_i$  and  $z_j$ , respectively. The model trains to bring together views from the same image and moves away views of different images, in that it minimizes  $\sum_{i=1}^N \ell_{i,j=pos(i)} + \sum_{j=1}^N \ell_{j,i=pos(j)}$ , with :

$$\ell_{i,j} = -\log \frac{\exp(\text{sim}(z_i, z_j)/\tau)}{\sum_{k=1, k \neq i}^{2N} \exp(\text{sim}(z_i, z_k)/\tau)} \quad (1)$$

$pos(i)$  represents the positive pair associated with  $i$ ,  $\tau$  is a temperature parameter, and  $\text{sim}(\dots)$  is the cosine similarity function.

## 2.5 Topology and graph-based augmentations

We test two different sets of topology-based transformations to generate the SimCLR views (see §2.8 for details on the parameters optimization and the right panel of Fig 3 for illustration). The first one, called "cutout", produces two different views: in one view, we remove a rolling block whose dimensions are 55% of the image volume dimensions. In the second view, all non-zero voxels inside a random 3D block of the same size are kept, and all non-zero voxels outside this block are removed.<sup>8</sup> In all views, fold bottom pixels are conserved to keep the information on the global longitudinal shape of the sulcus.

The second augmentation method is called "branch-clipping": for both views, it removes random branches from the folding graph until there are at least 40% of the voxels removed. All bottom voxels are removed.

We then binarize the image for all views and apply a random rotation with nearest-neighbor interpolation with a maximum amplitude of 6° around each axis.

## 2.6 Backbones

We tested three different backbones. The first is a simple six-layer convolutional network, taken from a  $\beta$ -VAE encoder optimized for cortical skeleton images.<sup>8</sup> The second one is a DenseNet architecture,<sup>22</sup> which gives good results when used on brain MRI data.<sup>23</sup> Its number of dense blocks is reduced to two, considering the smaller image size.<sup>8</sup>

The third is a PointNet.<sup>24</sup> A PointNet is a neural network that, instead of taking a binary 3D matrix with a fixed size as input, takes point clouds, i.e., a list of 3D coordinates corresponding to the non-zero pixels in the original 3D matrix. Indeed, in our volumes, only 4% of the pixels have non-zero values, making this backbone attractive. We take the original implementation of PointNet in this work.

## 2.7 Linear classification

A standard way to assess the quality of a SimCLR model is to train classifiers on the embeddings it generates.<sup>7</sup> We are using the scikit-learn implementation of a linear Support Vector Classifier (SVC) with probability estimates: we take as a criterion the area under the receiver operating characteristic curve (AUC).<sup>25</sup>

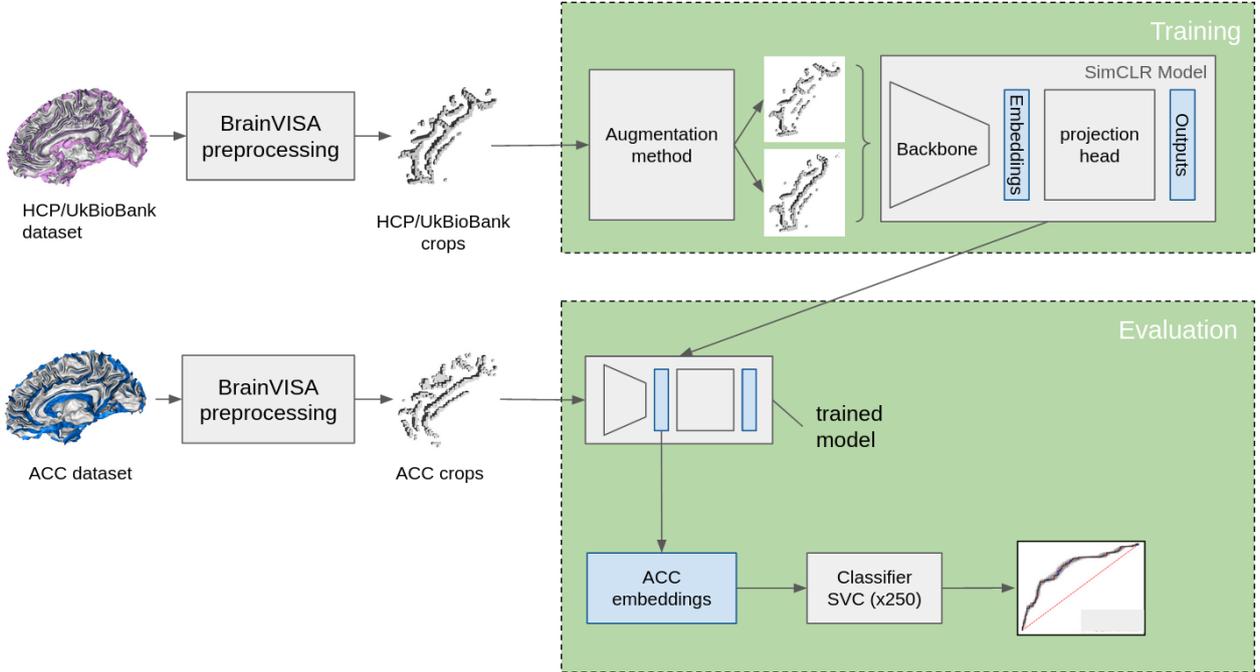


Figure 4. *Preprocessing and training pipeline.* We train the SimCLR model either on the HCP or on the UKBioBank dataset (top) and evaluate the model on the ACC dataset, which has been manually labeled for the double-parallel pattern in the cingulate region (bottom).

## 2.8 Deep-learning engineering and parameter optimization

We do initial grid searches using the DenseNet architecture with a 30-dimensional latent space and a non-linear projection head. We are first doing a gridsearch on HCP-1 with evaluation on ACC-1 with a 55% cutout augmentation<sup>8</sup> over the learning rate on 2,  $4 \times 10^{-4}$ , the batch size on 16, 32. We then set the learning rate to  $4 \times 10^{-4}$  and the batch size to 16. We then run a second grid search on 30,45,55% cutout size for the cutout augmentation strategy and on 30,40,50% of removed pixels for the branch-clipping strategy. We find 55% for the cutout size and 40% for the percentage of removed pixels to be the best.

We then test a dropout rate of 0, 0.05, 0.1 for the DenseNet backbone with both a 4-dimension and 30-dimension latent space, with both the cutout and the branch-clipping augmentations. We set the dropout rate to 0.05.

Last, we test the DenseNet backbone when removing or keeping the bottom voxels. For the cutout augmentation, keeping bottom voxels increases the probability of getting stuck in a trivial solution where all outputs lie in the same place, particularly for latent spaces of size 4. For the branch-clipping augmentation, removing bottom voxels improves the AUC, probably by increasing the randomization of the views. Thus, we kept all bottom voxels in the cutout augmentation and removed all bottom voxels in the branch-clipping augmentation.

We then run 5 models for each condition (see §3.1). Running a ConvNet model takes 20 minutes on a 16GB GPU (Quadro RT 5000) with 48 CPU cores.

## 3. EXPERIMENTS AND RESULTS

### 3.1 Optimizing SimCLR model for the detection of the double-parallel pattern

We optimize the SimCLR self-supervised model to detect the double-parallel pattern in the cingulate region. This is done in three main steps: we first perform a preprocessing of the brain MRI data, we then perform

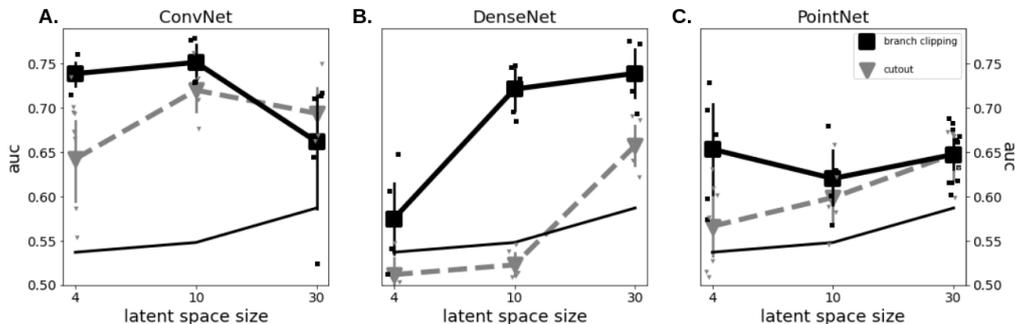


Figure 5. *Parameter optimization results.* Training is done on HCP-1 (half of the HCP dataset). Evaluation is done on ACC-1 (half of the ACC dataset). Each small marker represents a different trained model. **A.** AUC score as a function of the latent space size for the ConvNet backbone. Black lines with squares and dashed lines with down triangles stand respectively for the branch-clipping and the cutout augmentations. The thinner solid black line corresponds to a Principal Component Analysis (PCA) model. **B.** and **C.** are similar plots, respectively, for the DenseNet and the PointNet backbones.

parameter optimization on the train/validation dataset, and finally, we test and evaluate the model on the test dataset (Fig. 4).

After preprocessing the brain MRI data, we obtain crops of the cingulate region of the right hemisphere containing only information on folds (see §2.3).

We then optimize parameters by training the model on HCP-1 and evaluating it using a linear SVC on ACC-1 (see §2.8 for details). The evaluation criterion is the AUC for linearly classifying the double-parallel pattern in the cingulate region. After an initial parameter initialization, we test three backbones—a simple six-layer convolutional network (named ConvNet), a DenseNet, and a PointNet—, three latent spaces sizes—4, 10, 30—, two augmentation strategies—the cut-out and the branch-clipping augmentation—, and two projection heads—linear and non-linear—. We find the best model to be the six-layer convolutional network with a linear projection head, a latent space size (that is, the dimension of the latent space) of 10, and using the branch-clipping augmentation strategy (Fig. 5).

Last, for testing, we train models with the chosen hyperparameters on HCP-2 and evaluate them on ACC-2. We find the AUC of SimCLR on the latent space to be  $0.73 \pm 0.03$  ( $n=5$  models), slightly above the AUC ( $0.69 \pm 0.04$ ,  $n=5$ ) of the  $\beta$ -VAE model optimized in<sup>8</sup> and applied here to the same dataset (Fig. 6A). Even if favorable, we note that the difference doesn’t reach statistical significance ( $p=0.08$ ).

We observe that the effect of the batch size on the AUC reaches a plateau for a small batch size of 8 (Fig. 6B); we use a batch size of 16 in our models after an initial optimization (see §2.8). This differs from applying SimCLR on big 2D image datasets, in which bigger batch sizes improve the results.<sup>7</sup> Then, we observe that, by training on the ACC-1 dataset whose size is  $N=171$ , the result is similar to the one obtained when training the model on HCP-2 with a similar training size ( $p=0.84$ , t-test, comparing blue circle and black square in Fig. 6B). The result is also true for the UKBioBank dataset for enough training data, as the UKBioBank results match the HCP ones for 1101 subjects. As data acquisition sites and age ranges differ in all three datasets, this shows that the method is age and site-resistant. Then, when changing the training set size  $N$ , the performances increase. For UKBioBank, the AUC increases to 0.76 when training with 21070 subjects (see Fig. 6C).

Variability (measured as the standard deviation over 5 different SimCLR initializations) decreases from 0.04 for  $n=551$  to 0.01 for  $n=21070$ , showing that increasing the number of training subjects significantly stabilizes the representation. Last, the UMAP representation (from a model trained on UKBioBank and fit with UMAP on UKBioBank embeddings) shows how the ACC dataset, as well as the subjects with a double-parallel pattern, is embedded in the UKBioBank space (Fig. 6D).

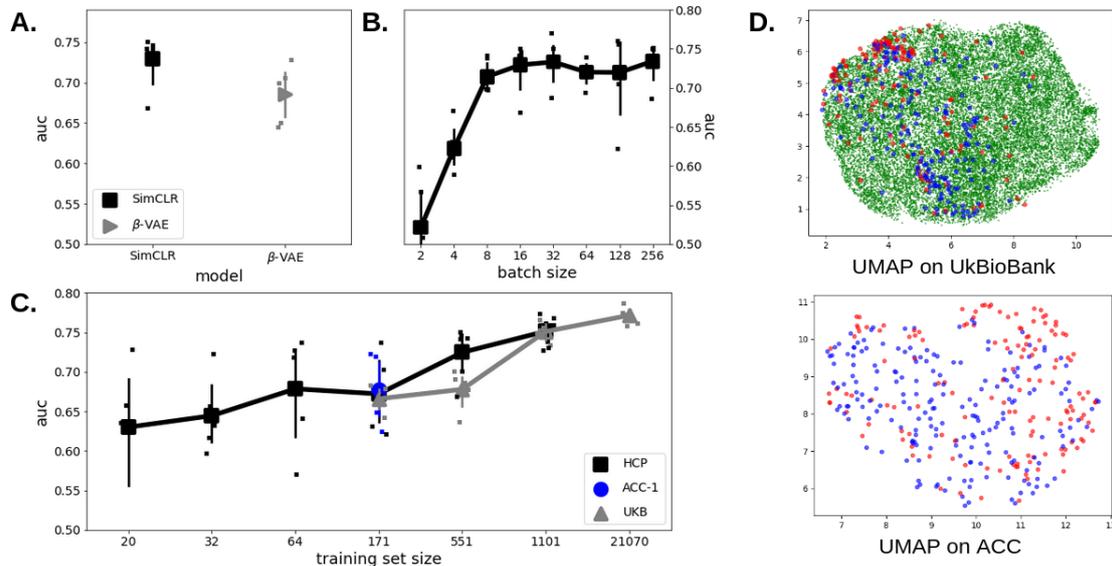


Figure 6. *Testing results with the best parameters for SimCLR.* Training is done on HCP-2 and evaluation is done on ACC-2. **A.** AUC for SimCLR compared with the  $\beta$ -VAE model. **B.** AUC as a function of batch size. **C.** AUC as a function of the training set size. The blue circle and the upper triangle represent models for which the training is done, respectively on ACC-1 and UKBioBank. HCP training is done exclusively on HCP-2 when the training size is smaller than 551. **D.** UMAP representations of UKBioBank embeddings (top) and of ACC embeddings (bottom). Each point represents a UKBioBank subject (green), an ACC subject with a double-parallel sulcus (red) or without (blue.)

#### 4. DISCUSSION AND CONCLUSIONS

We built the software structure and applied for the first time a self-supervised deep-learning algorithm on cortical skeletons with a database of more than 20,000 subjects. For the first time, we implemented augmentations that use the graph structure of the cortical folds. We optimized the model and quantified its significance using a manually labeled dataset.

We obtained a latent representation of the cortical folds of the cingulate region and evaluated this representation for detecting a double parallel cingulate pattern known to be related to schizophrenia.

The best SimCLR model found has two characteristics: a very simple backbone—a convolutional neural network—and a topology-specific augmentation—the branch-clipping—that differs from the augmentations used in 2D image learning and takes into account topological information of the inputs. Its simplicity likely means the same model will work reasonably well when applied to other brain regions.

The obtained representation quality is resilient to a change in the training database. This advantage likely comes from using cortical skeletons as they are less sensitive to acquisition sites and the subject’s age.<sup>15</sup> However, there is still a dataset effect. A possible solution would be to apply debiasing techniques<sup>26</sup> or to modify the augmentations further to reduce the remaining site and age effect.

Now that we have a representation that has been optimized for folding pattern detection, the next step will be to bring external information (akin to other modalities) directly into the model to further improve its representation quality.<sup>23</sup> A possible solution may be to perform a late fusion (that is, to concatenate the modalities at the level of the latent space), followed by an autoencoder to leverage possible correlations.<sup>27</sup>

#### Acknowledgments

This work was funded by the grants ANR-19-CE45-0022-01FOPASUBA and ANR-20-CHIA-0027-01FOLDDICO.

## REFERENCES

- [1] Mellerio, C., Roca, P., Chassoux, F., Danière, F., Cachia, A., Lion, S., Naggara, O., Devaux, B., Meder, J.-F., and Oppenheim, C., “The Power Button Sign: A Newly Described Central Sulcal Pattern on Surface Rendering MR Images of Type 2 Focal Cortical Dysplasia,” *Radiology* **274**, 500–507 (Feb. 2015).
- [2] Cachia, A., Borst, G., Vidal, J., Fischer, C., Pineau, A., Mangin, J.-F., and Houdé, O., “The Shape of the ACC Contributes to Cognitive Control Efficiency in Preschoolers,” *Journal of Cognitive Neuroscience* **26**, 96–106 (Jan. 2014).
- [3] Yucel, M., Stuart, G. W., Maruff, P., Velakoulis, D., Crowe, S. F., Savage, G., and Pantelis, C., “Hemispheric and Gender-related Differences in the Gross Morphology of the Anterior Cingulate/Paracingulate Cortex in Normal Volunteers: An MRI Morphometric Study,” *Cerebral Cortex* **11**, 17–25 (Jan. 2001).
- [4] Garrison, J. R., Fernyhough, C., McCarthy-Jones, S., Haggard, M., and Simons, J. S., “Paracingulate sulcus morphology is associated with hallucinations in the human brain,” *Nature Communications* **6**, 8956 (Nov. 2015).
- [5] Borne, L., Rivière, D., Cachia, A., Roca, P., Mellerio, C., Oppenheim, C., and Mangin, J.-F., “Automatic recognition of specific local cortical folding patterns,” *NeuroImage* **238**, 118208 (Sept. 2021).
- [6] Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A., “beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework,” *ICLR* (2017).
- [7] Chen, T., Kornblith, S., Norouzi, M., and Hinton, G., “A Simple Framework for Contrastive Learning of Visual Representations,” in [*Proceedings of the 37th International Conference on Machine Learning*], 1597–1607, PMLR (Nov. 2020).
- [8] Chavas, J., Guillon, L., Pascucci, M., Dufumier, B., Rivière, D., and Mangin, J.-F., “Unsupervised Representation Learning of Cingulate Cortical Folding Patterns,” in [*Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*], Wang, L., Dou, Q., Fletcher, P. T., Speidel, S., and Li, S., eds., *Lecture Notes in Computer Science*, 77–87, Springer Nature Switzerland, Cham (2022).
- [9] Guillon, L., Cagna, B., Dufumier, B., Chavas, J., Rivière, D., and Mangin, J.-F., “Detection of Abnormal Folding Patterns with Unsupervised Deep Generative Models,” in [*Machine Learning in Clinical Neuroimaging*], Abdulkadir, A., Kia, S. M., Habes, M., Kumar, V., Rondina, J. M., Tax, C., and Wolfers, T., eds., *Lecture Notes in Computer Science*, 63–72, Springer International Publishing, Cham (2021).
- [10] Guillon, L., Chavas, J., Bénézit, A., Moutard, M.-L., Roca, P., Mellerio, C., Oppenheim, C., Rivière, D., and Mangin, J.-F., “Identification of Rare Cortical Folding Patterns using Unsupervised Deep Learning,” *Imaging Neuroscience* (Jan. 2024).
- [11] Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., and Ugurbil, K., “The WU-Minn Human Connectome Project: An overview,” *NeuroImage* **80**, 62–79 (Oct. 2013).
- [12] Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., Liu, B., Matthews, P., Ong, G., Pell, J., Silman, A., Young, A., Sprosen, T., Peakman, T., and Collins, R., “UK Biobank: An Open Access Resource for Identifying the Causes of a Wide Range of Complex Diseases of Middle and Old Age,” *PLOS Medicine* **12**, e1001779 (Mar. 2015).
- [13] Chakravarty, M. M., Rapoport, J. L., Giedd, J. N., Raznahan, A., Shaw, P., Collins, D. L., Lerch, J. P., and Gogtay, N., “Striatal shape abnormalities as novel neurodevelopmental endophenotypes in schizophrenia: A longitudinal study,” *Human Brain Mapping* **36**(4), 1458–1469 (2015).
- [14] Rapoport, J. L. and Gogtay, N., “Childhood onset schizophrenia: support for a progressive neurodevelopmental disorder,” *International Journal of Developmental Neuroscience* **29**, 251–258 (May 2011).
- [15] Cachia, A., Borst, G., Tissier, C., Fisher, C., Plaze, M., Gay, O., Rivière, D., Gogtay, N., Giedd, J., Mangin, J.-F., Houdé, O., and Raznahan, A., “Longitudinal stability of the folding pattern of the anterior cingulate cortex during development,” *Developmental Cognitive Neuroscience* **19**, 122–127 (June 2016).
- [16] Delalande, L., Moyon, M., Tissier, C., Dorriere, V., Guillois, B., Mevell, K., Charron, S., Salvia, E., Poirel, N., Vidal, J., Lion, S., Oppenheim, C., Houdé, O., Cachia, A., and Borst, G., “Complex and subtle structural changes in prefrontal cortex induced by inhibitory control training from childhood to adolescence,” *Developmental Science* **23**(4), e12898 (2020).

- [17] Tissier, C., Linzarini, A., Allaire-Duquette, G., Mevel, K., Poirel, N., Dollfus, S., Etard, O., Orliac, F., Peyrin, C., Charron, S., Raznahan, A., Houdé, O., Borst, G., and Cachia, A., “Sulcal Polymorphisms of the IFC and ACC Contribute to Inhibitory Control Variability in Children and Adults,” *eNeuro* **5**, ENEURO.0197–17.2018 (Mar. 2018).
- [18] Ono, M., Kubik, S., and Abernathy, C. D., [*Atlas of the cerebral sulci*], G. Thieme Verlag ; Thieme Medical Publishers, Stuttgart; New York (1990). OCLC: 645306373.
- [19] Mangin, J.-F., Frouin, V., Bloch, I., Régis, J., and Lopez-Krahe, J., “From 3D magnetic resonance images to structural representations of the cortex topography using topology preserving deformations,” *Journal of Mathematical Imaging and Vision* **5**, 297–318 (Dec. 1995).
- [20] Rivière, D., Mangin, J.-F., Papadopoulos-Orfanos, D., Martinez, J.-M., Frouin, V., and Régis, J., “Automatic recognition of cortical sulci of the human brain using a congregation of neural networks,” *Medical Image Analysis* **6**, 77–92 (June 2002).
- [21] Borne, L., Rivière, D., Mancip, M., and Mangin, J.-F., “Automatic labeling of cortical sulci using patch- or CNN-based segmentation techniques combined with bottom-up geometric constraints,” *Medical Image Analysis* **62**, 101651 (May 2020).
- [22] Huang, G., Liu, Z., van der Maaten, L., and Weinberger, K. Q., “Densely Connected Convolutional Networks,” *arXiv:1608.06993 [cs]* (Jan. 2018). arXiv: 1608.06993.
- [23] Dufumier, B., Gori, P., Victor, J., Grigis, A., Wessa, M., Brambilla, P., Favre, P., Polosan, M., McDonald, C., Pigué, C. M., and Duchesnay, E., “Contrastive Learning with Continuous Proxy Meta-Data for 3D MRI Classification,” *arXiv:2106.08808 [cs, stat]* (June 2021). arXiv: 2106.08808.
- [24] Charles, R. Q., Su, H., Kaichun, M., and Guibas, L. J., “PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation,” in [*2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*], 77–85, IEEE, Honolulu, HI (July 2017).
- [25] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E., “Scikit-learn: Machine Learning in Python,” *Journal of Machine Learning Research* **12**(85), 2825–2830 (2011).
- [26] Tartaglione, E., Barbano, C. A., and Grangetto, M., “EnD: Entangling and Disentangling deep representations for bias correction,” in [*2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*], 13503–13512 (June 2021).
- [27] D’Souza, N. S., Wang, H., Giovannini, A., Foncubierta-Rodríguez, A., Beck, K. L., Boyko, O., and Syeda-Mahmood, T., “Fusing Modalities by Multiplexed Graph Neural Networks for Outcome Prediction in Tuberculosis,” in [*MICCAI 2022*], **13437**, 287–297 (2022).