Towards a foundation model for cortical folding

Julien Laval, Joël Chavas, Vanessa Troiani, William Snyder, Marisa Patti, Mylène Moyal, Marion Plaze, Arnaud Cachia, Zhong Yi Sun, Vincent Frouin, et al.

Towards a foundation model for cortical folding

Julien Laval[1], Joël Chavas[1], Vanessa Troiani[2], William Snyder[3,4], Marisa Patti[5], Mylène Moyal[6], Marion Plaze[6], Arnaud Cachia[7], Zhong Yi Sun[8], Vincent Frouin[1], Pietro Gori[9], Denis Rivière[1], Jean-François Mangin[1]

[1] Université Paris-Saclay, CEA, CNRS, NeuroSpin, U9027 Baobab, Saclay, France
[2] Geisinger Autism and Developmental Medicine Institute, Lewisburg, USA
[3] National Institute of Mental Health Intramural Research Program, Bethesda, USA
[4] Department of Psychiatry, University of Cambridge, Cambridge, UK
[5] A.J. Drexel Autism Institute, Drexel University, USA
[6] GHU Paris, IPNP, INSERM U1266, Paris, France
[7] Université de Paris, LaPsyDÉ, CNRS, Paris, France
[8] US52-UAR2031 CATI, Institut du Cerveau (ICM), Paris, France
[9] LTCI, Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France

Abstract. The brain surface is composed of humps called gyri, separated by grooves called sulci. Although the main folds are common to all individuals, their shape varies, making them unique to each individual. Cortical folding may contain biomarkers that have yet to be deciphered. While conventional geometric approaches fail to fully characterize the high inter-individual variability, recent efforts in large-scale MRI data collection allow us to leverage the statistical power of deep neural networks. Here, we introduce Champollion V0, a self-supervised learning (SSL) algorithm to sort sulcal variability based on 21,070 subjects from the UKBioBank dataset. We revisit from scratch an existing model and optimize its ability to retrieve hand-labeled patterns defined by the neuroscientific community. Under linear evaluation on the latent space, Champollion V0 significantly improves the detection of three different kinds of folding patterns: the presence of a parallel sulcus (AUC increases from 73% to 84%), the presence of specific interruptions (AUC increases from 50% to 79%) and the detection of a specific folding shape (R^2 increases on each of the six main geometric features), respectively in the cingulate, the orbital and the central region. These hand-labeled patterns were found to be correlated to neurodevelopmental pathologies. Champollion V0 could enable the automatic labeling of larger datasets for future studies. The code can be found on Github.

Keywords: Brain · folding patterns · MRI · Self-supervised learning.

1 Introduction

The brain is folded, formed of bumps (gyri) separated by folds (sulci). The major folds are common to all individuals. However, they vary in shape, making each brain unique. Some sulcal patterns have been identified as biomarkers : for example, the cingulate region can display either a single main sulcus, or

an additional paracingulate sulcus (PCS), parallel to the cingulate (Fig. 1a). The asymmetry of this pattern between the two hemispheres has been linked to control efficiency in preschoolers [4]. Similarly, different prevalences of the orbitofrontal cortex (OFC) interruption types (Fig. 1b) have been reported in schizophrenic and catatonic subjects compared to controls [16, 18].

The brain folds may contain many biologically relevant patterns that have yet to be discovered. However, identifying and labeling such patterns is very time-consuming, if not impossible for human experts. The recent acquisitions of large MRI datasets have enabled the use of deep learning to sort the sulcal variability and extract patterns automatically. Gaudin et al. developed a contrastive self-supervised (SSL) model (named thereafter the Orig. model) to encode region-specific sulcal variability [13]. They optimized it on the cingulate region, training on 551 subjects from the Human Connectome Project (HCP), with the PCS detection as downstream task. The model's generalizability allowed the detection of a right superior temporal sulcus pattern linked to extreme prematurity [17]. Likewise, using a supervised counterpart of the model, Chavas et al. identified regions correlated to inhibitory control [6].

However, Gaudin et al. did not observe a significant improvement when training on 21,070 subjects from UkBioBank [2] instead of HCP, which calls for a more thorough optimization on UkBioBank to leverage the dataset size. Moreover, their model has been optimized solely for PCS detection. Here, we speculate that this was due to augmentations whose semantic content was not rich enough and to a downstream task that was too simple for optimization. Therefore, we propose the following contributions: by adding as a downstream task a new dataset hand-labeled for complex folding patterns (the orbitofrontal cortex patterns) and by using UkBioBank for training, we develop a new specific augmentation, TrimDepth, and revisit the augmentation pool, the loss, and the backbone. We establish a new state-of-the-art model and prove its ability to detect parallel sulci, their interruptions, and their shape. We name it Champollion V0, after the man who deciphered hieroglyphics, since the present work aims to decipher the hitherto unknown language of cortical folding.

2 Methods

Input data : Structural MR images of the brain are processed through the BrainVisa Morphologist pipeline[10] that produces a skeletonized negative cast of the brain. It transforms the sulci into surfaces (3D objects of one-voxel width) following the middle of the folding perpendicularly to the brain hull. These cortical skeletons are affinely normalized in the Talairach space with a 2mm voxel size while constrained to keep a one-voxel width. This preprocessing is meant to emphasize the folding while reducing the bias caused by the original resolution of the acquisition site. Then, a region of interest (ROI) is cropped, using the deep_folding toolbox [7] [11], to focus on specific sulci (same tool as

[10] https://brainvisa.info
[11] github: neurospin/deep_folding

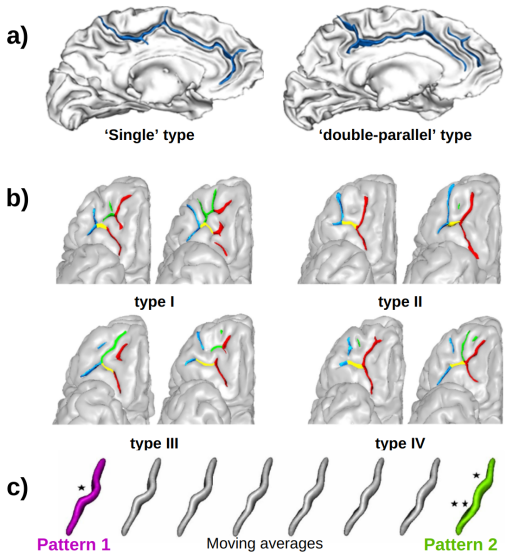

Fig. 1. The folding patterns studied in this paper. **a)** The two cingulate patterns (blue): the "single" type, with only the cingulate sulcus, and the "double parallel" type, with an additional paracingulate sulcus (PCS) (courtesy from [4]). **b)** Four orbitofrontal cortex pattern types with individual labeled orbitofrontal sulci, from the HCP database. Medial orbital sulcus (MOS) is labeled in red, lateral orbital sulcus (LOS) in blue, intermediate orbital sulcus (IOS) in green, and transverse orbital sulcus (TOS) in yellow. The Type I pattern has a discontinuous MOS and continuous LOS, the Type II a continuous MOS and continuous LOS, the Type III a discontinuous MOS and discontinuous LOS, and Type IV a continuous MOS and discontinuous LOS (courtesy from [24]). **c)**. Moving average of the Isomap first dimension on the central sulcus reveals a continuous change from a single knob (purple) to a double knob (green) pattern (courtesy from [20]).

Gaudin et al.). Since all datasets are normalized in the same space, a single dataset with hand-labeled sulci is required to define the mask of the ROI. We use a custom dataset (n=62) independent from the datasets used in this study to define the crops of the cingulate, orbitofrontal, and central regions.

Datasets and patterns description : The three datasets used in this paper are the following :

- **UkBioBank** [2] (n=21070) is a general population cohort used for SSL training.
- **ACCpatterns** (n=341) is a dataset composed of subjects taken from [5, 19, 3, 9, 23], with a hand-labeled paracingulate sulcus (PCS, Fig. 1a), our first studied pattern. It is used by Gaudin et al. to optimize the Orig. model, and like Gaudin et al., we focus on the right hemisphere.
- **Human Connectome Project** [26] (HCP). A subset (n=577) was hand-labeled in the orbitofrontal cortex (OFC) according to the four main inter-

ruption patterns (Fig. 1b) [24]. This is our second investigated pattern. We focus on the left hemisphere because the patterns are more uniformly distributed than in the right hemisphere [24]. Their distribution in our data is the following : Type I 49%, Type II 28%, Type III 17%, Type IV 6%. These subjects were also assigned six continuous shape descriptors for their central sulcus using Isomap [22], on a shape similarity matrix [21]. The associated regression tasks are not used during optimization, but only for final evaluation. The first Isomap dimension of the central sulcus is shown in Fig. 1c.

For ACCpatterns and HCP, the data was split into 10 stratified folds according to the label (PCS for ACCpatterns and OFC for HCP), sex, and acquisition sites. Moreover, HCP siblings were systematically assigned to the same splits. The SSL optimization of each task was conducted via 8-fold cross-validation on 8 out of the 10 stratified splits. The remaining 2 splits were retained for the final evaluation. In the SSL setting, ACCpatterns and HCP were only used for the downstream tasks. In the supervised baseline settings (described below), the same 2 splits were utilized for testing, 7 for training and 1 for validation.

Self-supervised learning principle and losses : For each image in a batch, two views are generated using random data augmentations, forming a positive pair. SSL losses bring together the positive pairs in the latent space, ensuring invariance to the augmentations to build semantically expressive representations. To avoid trivial solutions, contrastive methods such as SimCLR consider two views originating from different images as a negative pair and push them away in the latent space [8]. Conversely, BarlowTwins reduces redundancy by decorrelating the latent variables [27]. In this paper, we compare SimCLR, used in the Orig. model, to BarlowTwins.

Although SimCLR is known to perform better with large batch sizes on natural images [8], Gaudin et al. observed a plateau beyond a batch size of 16 [13]. We find the same results when training on a much larger database and hypothesize that the number of latent classes being small, negative pairs in large batches are likely to belong to the same class. Conversely, BarlowTwins is considered a negative-sample-free method [25] and thus looks more suitable.

Hyperparameters : For SimCLR, we use a temperature $\tau = 0.1$ and a batch size of 16 as in the Orig. model. BarlowTwins is reportedly not sensitive to batch size [27]. We set it to 32 as done in [25] with a similar dataset size. According to [27], the regularization hyperparameter λ should be close to $\frac{1}{d}, d$ being the latent space size. $\lambda = 10^{-2} (= \frac{2.56}{d}$ with $d = 256)$ is selected after search in the range $(5 * 10^{-3} - 10^{-2})$. The other hyperparameters are taken from the Orig. model: 250 epochs, a learning rate of $4 * 10^{-4}$, and a dropout rate of 5% in the backbone.

Domain-specific augmentations : The efficacy of SSL is contingent upon the meticulous design of augmentations [15, 8, 1] that must maintain the semantics while introducing sufficient variation within the input space. However, there is currently no consensus on defining good augmentations for medical imaging due to the lack of prior knowledge of the semantic content. Consequently, this topic is often overlooked [11].

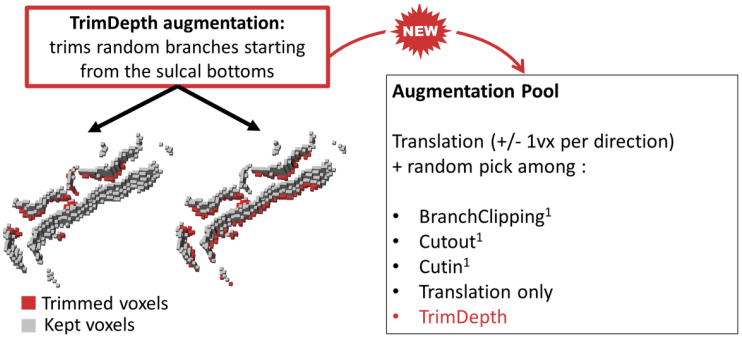

Fig. 2. (Left) Two views examples using TrimDepth. (Right) The augmentation pool. The augmentations marked by [1] have been first defined in [13]

.

In the context of cortical skeletons, typical natural image augmentations such as flips or color jittering are not adapted, given that the data is normalized in space and binary. In [13], Gaudin et al. craft domain-specific augmentations for cortical skeletons. Inspired by the popular Cutout [10], they design BranchClipping, which masks skeleton branches until a given percentage of the positive voxels are removed. They also try Cutout and Cutin, which keeps the inside of the mask instead of the outside, but they eventually select BranchClipping as the main augmentation and combine it with a small rotation.

In this paper, we introduce a new augmentation, TrimDepth, which randomly selects folds in the image to trim them of a given depth starting from the bottom. This should help the algorithm focus on sulci shapes rather than depth, which is more relevant. Indeed, sulci shapes are fixed throughout life once the brain is folded but sulci become more shallow with aging. In [3], Cachia et al. define folding patterns as objects invariant with age. Therefore, this augmentation should preserve the subjects' sulcal signatures.

Furthermore, we replace the rotations with small translations (they are faster to compute and don't risk affecting the semantics of images that are not rotation invariant). Differently from Gaudin et al., which look for the augmentation that gives the best result, we hypothesize that using a diverse set of augmentations prevents from learning their individual biases (*e.g.* the PCS may be erased by BranchClipping). We here redesign the augmentation framework so that one augmentation among the following pool is picked for each view with equal probability (except for the translation, which is always applied): BranchClipping, Cutout, Cutin, TrimDepth. By doing a gridsearch, we find that combining all the augmentations yields the best results. Because all the augmentations remove voxels, we do not allow them to be mixed on the same view to prevent the erasure of too many voxels. Additionally, with a probability of 0.2, we apply solely the translation so that the algorithm can learn to represent full images. For Cutout/Cutin, we select a mask covering 30% of the voxels and get the best results when keeping the bottom lines of the masked folds. We hypothesize that

maintaining a trace of the global topology helps to rely on the entire image, as opposed to natural images where the semantics are primarily local (*e.g.* an object on a background). For BranchClipping, we find that removing too many branches is detrimental to the OFC classification, so we choose to remove a single branch. But as it may remove a very small number of voxels in some cases, we remove all the bottom voxels aswell. Finally, we select a 2mm depth to be removed by TrimDepth. An illustration of TrimDepth and the new augmentation pool is presented in Fig. 2.

Backbones : We start with the Orig. model's six-layer convolutional network and examine deeper architectures. First, we double the number of layers and filters and increase the initial kernel size from 3 to 7 for 2mm resolution inputs and to 11 for 1.5mm resolution. Second, we try a deeper network, the 3D ResNet18 [14] (implemented by [11]). To save computation time, the models using a larger backbone than the original one are trained for 70 epochs, which is found to be enough for convergence.

Model Evaluation : We use the Area Under the Receiver Operating Characteristic Curve (AUC) for classification and the coefficient of determination R^2 for regression. In the multiclass case, we measure the AUCs in a One-vs-Rest scheme and report their weighted average. Deep learning models are trained 5 times to compute a standard deviation. SSL performance is evaluated using a linear model on the latent space (here a linear SVM), the standard way to assess representation quality.

Baselines : The SSL is compared to several baselines of increasing complexity:

- **PCA :** a Principal Component Analysis (PCA) is fit on the UkBioBank skeletons to reduce dimensionality to the latent space size, applied to the datasets of interest, and followed by a linear SVM.
- **Linear models** (Logistic regression, ElasticNet) and **rbf-SVM** : we use these supervised methods on the skeletons as they are classically used to compare with deep learning models in brain MRIs [11]. The regularization and other hyperparameters are detailed in Appendix Table 1.
- **Supervised deep learning :** we tried to either use the SSL augmentation pool or remove the augmentations except for the translation. We obtained best results without the augmentation pool for both classification tasks. We tested the best convolutional network found for SSL and a ResNet18 and found the 12-layer ConvNet backbone to be better for both tasks. We train for 200 epochs and use early stopping to limit overfitting. To handle the OFC class imbalance, we try weighting the loss, but don't achieve better results. The hyperparameter search is detailed in Appendix Table 2.

3 Experiments and Results

SSL optimization : The Orig. model is here completely redesigned, and the steps are illustrated in Fig. 3. First, we find that increasing the dimension of

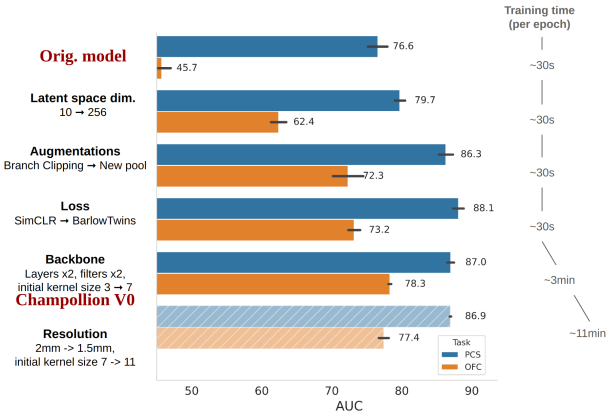

Fig. 3. Summary of the SSL improvements for the PCS and OFC classification. Each model is evaluated with linear SVM using 8-fold cross-validation on the train/validation splits. Training time is reported for a Quadro RTX 5000.

the latent space (from 10 to 256) is critical for capturing the pattern information in the OFC. Gaudin et al. selected a latent space size of 10 but did not explore values beyond 30, as the parameter did not appear to benefit the PCS detection. Then, the augmentations are replaced with the new pool (Fig. 2). We observe a dramatic increase in performance, especially on the OFC patterns, without additional computational cost. Next, we replace the SimCLR loss with BarlowTwins. It leads to a slight improvement in both tasks and better stability (a standard deviation of 2.5% for SimCLR against 1.0% for BarlowTwins on the OFC patterns), although it does not appear as crucial as the choice of the augmentations.

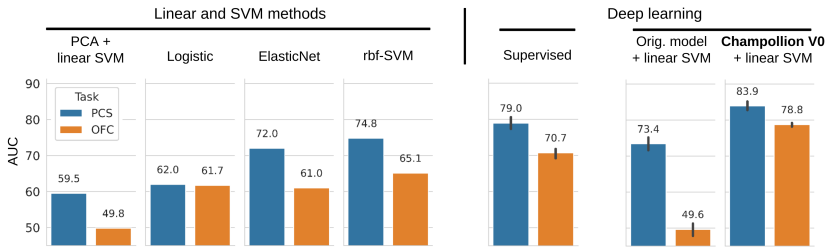

Fig. 4. Comparison between SSL and different baselines on the PCS detection and OFC patterns classification. Each model is evaluated on the test set.

Furthermore, we investigate larger backbones and achieve a 5% AUC improvement on the OFC patterns (at a non-negligible computational cost). First, we double the number of layers and filters of the original backbone. Subsequently,

we observe two distinct performance regimes while training multiple models, with an AUC for OFC detection of approximately 78% for some models and 65% for others, with no obvious difference in the loss value. We hypothesize that the first convolutional layer may be crucial for capturing interruptions before the deeper layers with stride potentially erase them, which may cause instability. Drawing inspiration from ResNet18, we increase the initial kernel size from 3 to 7, which stabilizes the learning, with all models giving an AUC of $\sim$ 78%. We try to further increase the backbone size using ResNet18 but we only achieve random performance. We hypothesize that it may be oversized for our problem.

Finally, we attempt to increase the image resolution to 1.5mm. We find that the initial kernel must be scaled to 11 for stability, but we do not observe any improvement, while training is significantly slowed down. At first, the results were surprising, as we expected the interruptions to be clearer at higher resolution, leading to a better classification of the patterns. However, a quick visualization indicates that the interruptions at 1.5mm are still visible at 2mm.

Champollion V0 outperforms the Orig. model on the test set by 29 points on OFC and 11 points on PCS (AUC, Fig. 4). Although our hyperparameter search may not have been exhaustive, we could not beat SSL with a supervised algorithm. Moreover, all supervised methods failed to classify the type IV OFC pattern (not shown on graph), present in only $\sim$ 6% of the population. By contrast, Champollion V0 performs well on all OFC types (type I: $77.2 \pm 0.6\%$, type II: $83.1 \pm 0.7\%$, type III: $77.4 \pm 0.3\%$, type IV: $81.4 \pm 1.4\%$). Thus, SSL pretraining appears to be particularly valuable for detecting rare patterns.

Application to Isomap regression : We compare Champollion V0 and the Orig. model in their ability to encode sulcus shape features, which was not assessed by the tasks used for optimization, by performing linear support vector regressions between the latent space of the central sulcus region and the first six dimensions of the central sulcus Isomap. The first dimension is illustrated in Fig 1c. Champollion V0 encodes all these geometrical features better (Table 1).

Isomap Dim.	1	2	3	4	5	6
Orig. Model	48.2 ± 2.2	42.1 ± 3.2	43.7 ± 12.3	34.5 ± 1.5	0.1 ± 0.1	11.3 ± 2.1
Champollion V0	$\mathbf{54.8} \pm 1.2$	$\mathbf{58.5} \pm 2.3$	$\mathbf{68.8} \pm 1.0$	$\mathbf{56.6} \pm 1.8$	$\mathbf{26.3} \pm 1.5$	$\mathbf{38.7} \pm 1.3$

Table 1. Comparison of the regression coefficient R^2 between the latent space and the 6 Isomap components, computed on the HCP test set.

4 Conclusion and perspectives

This paper describes Champollion V0, a region-specific encoder that provably detects sulci, their shape, and how they are interrupted. The algorithm achieves 84% AUC on PCS detection and 79% on OFC pattern classification. Given the

often ambiguous boundary between classes (*e.g.*, a very short PCS), it is likely that Champollion V0 has reached a plateau for these humanly defined patterns.

We underscore the design of domain-specific augmentations as a cornerstone of our model, while it is largely overlooked in medical imaging [11]. We believe always enriching the augmentation set is the path to further improvements. A possible augmentation would be a lateral trimming of the sulci to highlight the interruptions. Another improvement would be to combine the augmentations differently as large augmentation combinatorics benefit SSL [15].

To further establish Champollion as a foundation model for cortical folding, it will soon be trained on more subjects, 50,000 UkBioBank subjects. Second, we can enforce the latent space to encode biologically relevant patterns using supervision (*e.g.*, pathological patterns) and y-aware loss with metadata [12] (*e.g.*, genetically driven patterns using UkBioBank genetic distances).

We have shown that Champollion V0 bridges the gap between automated and manual labeling of folding patterns. It can now automatically detect known folding patterns for large-scale studies. This has a strong implication regarding the links between folding patterns and clinical endpoints. For example, the OFC patterns have been linked to schizophrenia, but only on small databases [16]. Champollion V0 can now permit the study of such links between folding patterns and schizophrenia on all available schizophrenia datasets.

Acknowledgments. This project has been funded by ANR via FOLDDICO (Project-ANR-20-CHIA-0027), BHT (ProjetIA-22-PESN-0012) and StratifyAging (ProjetIA-22-PESN-0010)

Disclosure of Interests. The authors have no competing interests to declare.

References

1. Balestriero, R., Ibrahim, M., Sobal, V., Morcos, A., Shekhar, S., Goldstein, T., Bordes, F., Bardes, A., Mialon, G., Tian, Y., et al.: A cookbook of self-supervised learning. arXiv preprint arXiv:2304.12210 (2023)
2. Bycroft, C., Freeman, C., Petkova, D., Band, G., Elliott, L.T., Sharp, K., Motyer, A., Vukcevic, D., Delaneau, O., O'Connell, J., Cortes, A., Welsh, S., Young, A., Effingham, M., McVean, G., Leslie, S., Allen, N., Donnelly, P., Marchini, J.: The UK Biobank resource with deep phenotyping and genomic data. Nature **562**(7726), 203–209 (Oct 2018), publisher: Nature Publishing Group
3. Cachia, A., Borst, G., Tissier, C., Fisher, C., Plaze, M., Gay, O., Rivière, D., Gogtay, N., Giedd, J., Mangin, J.F., Houdé, O., Raznahan, A.: Longitudinal stability of the folding pattern of the anterior cingulate cortex during development. Developmental Cognitive Neuroscience **19**, 122–127 (Jun 2016)
4. Cachia, A., Borst, G., Vidal, J., Fischer, C., Pineau, A., Mangin, J.F., Houdé, O.: The shape of the ACC contributes to cognitive control efficiency in preschoolers. Journal of Cognitive Neuroscience **26**(1), 96–106 (Jan 2014)
5. Chakravarty, M.M., Rapoport, J.L., Giedd, J.N., Raznahan, A., Shaw, P., Collins, D.L., Lerch, J.P., Gogtay, N.: Striatal shape abnormalities as novel neurodevelopmental endophenotypes in schizophrenia: A longitudinal study. Human Brain Mapping **36**(4), 1458–1469 (Dec 2014)
6. Chavas, J., Gaudin, A., Rivière, D., Mangin, J.F.: Regional supervised learning of inhibitory control strength from cortical sulci. In: Medical Imaging with Deep Learning (2024)
7. Chavas, J., Guillon, L., Pascucci, M., Dufumier, B., Rivière, D., Mangin, J.F.: Unsupervised Representation Learning of Cingulate Cortical Folding Patterns. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2022. pp. 77–87. Springer Nature Switzerland, Cham (2022)
8. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.E.: A simple framework for contrastive learning of visual representations. International conference on machine learning p. 1597–1607 (2020)
9. 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., Borst, G.: Complex and subtle structural changes in prefrontal cortex induced by inhibitory control training from childhood to adolescence. Developmental Science **23**(4), e12898 (Jul 2020)
10. DeVries, T., Taylor, G.W.: Improved Regularization of Convolutional Neural Networks with Cutout (Nov 2017), arXiv:1708.04552 [cs]
11. Dufumier, B., Gori, P., Petiton, S., Louiset, R., Mangin, J.F., Grigis, A., Duchesnay, E.: Exploring the potential of representation and transfer learning for anatomical neuroimaging: Application to psychiatry. NeuroImage **296**, 120665 (2024)
12. Dufumier, B., Gori, P., Victor, J., Grigis, A., Wessa, M., Brambilla, P., Favre, P., Polosan, M., McDonald, C., Piguet, C.M., Phillips, M., Eyler, L., Duchesnay, E.: Contrastive learning with continuous proxy meta-data for 3d mri classification. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2021. pp. 58–68. Springer International Publishing, Cham (2021)
13. Gaudin, A., Guillon, L., Fischer, C., Cachia, A., Rivière, D., Mangin, J.F., Chavas, J.: Optimizing contrastive learning for cortical folding pattern detection. In: Col-

liot, O., Mitra, J. (eds.) Medical Imaging 2024: Image Processing. vol. 12926, p. 129260Q. International Society for Optics and Photonics, SPIE (2024)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
15. Huang, W., Yi, M., Zhao, X., Jiang, Z.: Towards the Generalization of Contrastive Self-Supervised Learning (Mar 2023), arXiv:2111.00743 [cs, stat]
16. Isomura, S., Hashimoto, R., Nakamura, M., Hirano, Y., Yamashita, F., Jimbo, S., Yamamori, H., Fujimoto, M., Yasuda, Y., Mears, R.P., Onitsuka, T.: Altered sulcogyral patterns of orbitofrontal cortex in a large cohort of patients with schizophrenia. NPJ schizophrenia **3**, 3 (2017)
17. Laval, J., Rivière, D., Gaudin, A., Frouin, V., Dubois, J., Gondova, A., Mangin, J.F., Chavas, J.: Self-supervised contrastive learning unveils cortical folding pattern linked to prematurity. In: Medical Imaging with Deep Learning (2024)
18. Moyal, M., Haroche, A., Attali, D., Dadi, G., Raoelison, M., Berre, A.L., Iftimovici, A., Chaumette, B., Leroy, S., Charron, S., Debacker, C., Oppenheim, C., Cachia, A., Plaze, M.: Orbitofrontal sulcal patterns in catatonia. European Psychiatry **67**(1), e6 (Jan 2024)
19. Rapoport, J.L., Gogtay, N.: Childhood onset schizophrenia: support for a progressive neurodevelopmental disorder. International Journal of Developmental Neuroscience: The Official Journal of the International Society for Developmental Neuroscience **29**(3), 251–258 (May 2011)
20. Sun, Z.Y., Pinel, P., Rivière, D., Moreno, A., Dehaene, S., Mangin, J.F.: Linking morphological and functional variability in hand movement and silent reading. Brain Structure & Function **221**(7), 3361–3371 (Sep 2016)
21. Sun, Z.Y., Klöppel, S., Rivière, D., Perrot, M., Frackowiak, R., Siebner, H., Mangin, J.F.: The effect of handedness on the shape of the central sulcus. NeuroImage **60**(1), 332–339 (Mar 2012)
22. Tenenbaum, J.B., Silva, V.D., Langford, J.C.: A Global Geometric Framework for Nonlinear Dimensionality Reduction. Science **290**(5500), 2319–2323 (Dec 2000)
23. 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., Cachia, A.: Sulcal Polymorphisms of the IFC and ACC Contribute to Inhibitory Control Variability in Children and Adults. eNeuro **5**(1), ENEURO.0197–17.2018 (2018)
24. Troiani, V., Snyder, W., Kozick, S., Patti, M.A., Beiler, D.: Variability and concordance of sulcal patterns in the orbitofrontal cortex: A twin study. Psychiatry Research: Neuroimaging **324**, 111492 (Aug 2022)
25. Tsai, Y.H.H., Bai, S., Morency, L.P., Salakhutdinov, R.: A Note on Connecting Barlow Twins with Negative-Sample-Free Contrastive Learning (Apr 2021), arXiv:2104.13712 [cs]
26. Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E.J., Yacoub, E., Ugurbil, K.: The WU-Minn Human Connectome Project: An overview. NeuroImage **80**, 62–79 (Oct 2013)
27. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. International Conference on Machine Learning p. 12310–12320 (2021)