



HAL
open science

Unsupervised Representation Learning of Cingulate Cortical Folding Patterns

Joel Chavas, Louise Guillon, Marco Pascucci, Benoit Dufumier, Denis Rivière,
Jean-François Mangin

► **To cite this version:**

Joel Chavas, Louise Guillon, Marco Pascucci, Benoit Dufumier, Denis Rivière, et al.. Unsupervised Representation Learning of Cingulate Cortical Folding Patterns. MICCAI 2022, Sep 2022, Singapore, Singapore. hal-03685330v1

HAL Id: hal-03685330

<https://hal.science/hal-03685330v1>

Submitted on 2 Jun 2022 (v1), last revised 12 Sep 2022 (v2)

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.

Unsupervised Representation Learning of Cingulate Cortical Folding Patterns

Joël Chavas^{*1[0000000235595919]}, Louise Guillon^{*1[0000000157160070]}, Marco Pascucci^{1[0000-0002-3150-256X]}, Benoit Dufumier^{1,2[0000000282532363]}, Denis Rivière^{1[0000000219022213]}, and Jean-François Mangin^{1[0000-0002-1612-461X]}

¹ Université Paris-Saclay, CEA, CNRS, NeuroSpin, Baobab, Gif-sur-Yvette, France

`louise.guillon@cea.fr` or `joel.chavas@cea.fr`

² LTCI, Télécom Paris, IPParis, Paris, France

Abstract. The human cerebral cortex is folded, making sulci and gyri over the whole cortical surface. Folding presents a very high inter-subject variability, and some neurodevelopmental disorders are correlated to local folding structures, named folding patterns. However, it is tough to characterize these patterns manually or semi-automatically using geometric distances. Here, we propose a new methodology to identify typical folding patterns. We focus on the cingulate region, known to have a clinical interest, using so-called skeletons (3D representation of folding patterns). We compare two models, β -VAE and SimCLR, in an unsupervised setting to learn a relevant representation of these patterns. We add a decoder to SimCLR to be able to analyse latent space. Specifically, we leverage the data augmentations used in SimCLR to propose a novel kind of augmentations based on folding topology. We then apply a clustering on the latent space. Cluster folding averages, interpolation in the latent space and reconstructions reveal new pattern structures. This structured representation shows that unsupervised learning can help in the discovery of still unknown patterns. We will gain further insights into folding patterns by using new priors in the unsupervised algorithms and integrating other brain data modalities. Code and experiments are available at github.com/neurospin-projects/2021_jchavas_lguillon_deepcingulate.

Keywords: beta-VAE, SimCLR, contrastive learning, folding pattern, cortex.

1 Introduction

The human cortex is convoluted, made of folds, called gyri, separated by grooves, the sulci. Contrary to macaque, whose cortical folding follows a systematic scheme, human cortex folding is highly variable, making it a fingerprint of each individual. Although this diversity seems, first, intractable, neuroanatomists have succeeded in defining a partially reproducible scheme, which has led to the nomenclature of sulci used in neuroscience [21]. But each sulcus can have

* Contributed equally

a large number of patterns, which hinders its reliable identification (Fig. 1B). Deep learning could be a real lever to deal with this tremendous inter-individual variability.

Shapes of the folding patterns are particularly interesting to study as they are "trait" features, they remain during lifespan, contrary to "state" features (e.g. sulci depth or width) that are not stable throughout life [3]. Some works tried to decipher folding patterns and identify the most common shapes. Historically, this was done visually [28], enabling to define central sulcus knob and the omega-shape of the mid-fusifiform sulcus in particular [29,27]. However, manually finding relevant geometrical shapes is very hard due to the high diversity of folding patterns. Thus some studies tried to automate the characterization of folding patterns mainly based on geometric distances [25,20,8]. More recently, [24] trained neural network classifiers to map geometric shapes to folding patterns applied to the broken-H shape pattern in the orbitofrontal region.

However, characterizing the full diversity of folding patterns remains out of reach for these automatic geometric methods. Unsupervised deep learning methods is a natural next step: they have been used for detecting anomalies in folding shapes [12], but they have not been used yet to characterize the normal inter-individual variability of folding patterns.

Numerous approaches try to tackle unsupervised representation learning problems. On the one hand, auto-encoders (AE) are generative models that build a latent space comprising much fewer dimensions than the input, suggesting that the representations could be more easily understood, leading eventually to pattern discovery. For example, [12] showed that β -VAE are promising to detect anomalies of folding patterns.

On the other hand, self-supervised methods, particularly contrastive learning models, have proved to be very powerful. The foundation contrastive model, SimCLR [6], permits structuring the obtained latent space without using any labels. Its strength lies in the possibility to integrate prior information either by choosing the adapted random augmentations or by integrating into the loss function similarity information from other modalities.

Many works start to apply such framework to biomedical imaging. Thus, [26] proposed 3D versions of several self-supervised tasks on various objectives including brain tumor segmentation. Self-supervised methods offer the opportunity to leverage additional prior information from medical data. For instance, [10] applied contrastive learning to brain MRI and took advantage of available meta-data such as age and sex. This accelerating research on nearby fields shows that it is the right moment to apply self-supervised learning to the folding pattern characterization problem.

This study aims to pave the way for unsupervised deep learning to systematize the identification of typical folding patterns across the cortex in the future. More specifically, we aim to compare two unsupervised deep learning models in the task of obtaining a latent space structured enough to bring out folding patterns. To achieve this goal, we developed a deep learning pipeline that focuses on the folding pattern of predefined regions. We tested the pipeline on

the cingulate region, as it is sufficiently variable to justify the use of our methods, and it has a clinical interest for psychiatric disorders [30,22,2]. Then, we chose, adapted and compared two powerful and standard unsupervised methods, namely a contrastive learning model, SimCLR [6] to which we added a decoder, and a generative model, β -VAE [13]. Last, we proposed ways to analyze the results which are new and challenging with respect to classical deep learning literature as the input sample topologies are very different from classical 3D images.

Our contributions are three-fold. First, we implement an efficient topology-based augmentation for SimCLR. Second, we propose models reconstruction of cingulate patterns as alternative to local average folding pattern to explain the encoded features and last, a preprocessing pipeline to study efficiently folding patterns. Our work constitutes a first step to increase the knowledge on folding patterns that will permit to find biomarkers for brain disease.

2 Methods

2.1 Pre-processing

From brain MRI images, we used skeletons whose concept was first introduced in [19]. They consist in 3D images of the cortical folds obtained with BrainVISA/Morphologist preprocessing pipeline (<https://brainvisa.info/web/>). Skeletons' voxels are divided into background and folds (Figure 1A). Fold voxels can hold several values depending on their topological meaning (fold bottom, fold junction, etc.). Using this input enables to focus on the folding geometry and eliminates some biases such as age or site.

We focus our study on the cingulate region of the right hemisphere (Fig. 1A). We learned a mask of the cingulate and paracingulate sulci over a database where the folds were manually labelled [1]. In short, labeled subjects were first affinely normalized to a standard brain referential (ICBMc2009); then, each subject voxel belonging to the sulci of interest increments a sulcus-specific mask. We combined and dilated these two resulting masks to get a simple Region of Interest (ROI). We then applied this final mask to skeleton images of any unlabeled brain. Our final input is a 2-mm resolution 3D crop of dimension 20x40x40 (Fig. 1B) with integer values representing local topologies.

2.2 Learning cingulate region representations

We compared two unsupervised deep learning models : an autoencoder-based model and a contrastive learning framework.

β -VAE. AE-based models are commonly used to learn representations and to model the inter-subject variability. With an encoder θ , they enable to project data from input space \mathcal{X} onto a latent space \mathcal{Z} comprising much fewer dimensions. The latent code is then reconstructed thanks to a decoder ϕ . β -VAE [13],




Fig. 1. Framework to study skeletons in the cingulate region of the HCP dataset. **A)** Sample crop of the cingulate region represented as buckets superimposed with the white matter mesh. Red voxels are *bottom voxels*. **B)** Samples of the studied crops, given as inputs to the unsupervised algorithms. **C)** Whole pipeline. We generate crops of the cingulate region based on a manually labeled dataset. We train both models (β -VAE and SimCLR), we then infer and perform downstream analysis of our two models.

an extension of VAE [15], is particularly interesting as the latent space is constrained to follow a prior distribution and input data are encoded as a distribution. The objective function is a combination of the reconstruction error and the matching of two distributions using the Kullback-Liebler (KL) divergence. β -VAE is trained to maximize:

$$\mathcal{L}(\theta, \phi; \mathbf{x}, \mathbf{z}, \beta) = E_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] - \beta \mathcal{D}_{KL}(q_{\phi}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) \quad (1)$$

where $p(\mathbf{z})$ corresponds to the prior distribution (here, a reduced centered Gaussian distribution) and $q_{\phi}(\mathbf{z}|\mathbf{x})$, the posterior distribution. We ran the model on binarized skeletons.

SimCLR. SimCLR is an instance discrimination contrastive model. For each sample x of the batch of size N , we generate at each epoch two views x_i and x_j , whose model outputs are respectively z_i and z_j . The model trains to bring together views from the same image, that is it minimizes $\sum_{i=1}^N \ell_{i,j=pos(i)} + \sum_{j=1}^N \ell_{j,i=pos(i)}$, $\ell_{i,j}$ being the loss function for a positive pair of examples (τ is a temperature parameter) :

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

View generations are the algorithms specific to our problem: they use the discrete topology of the fold skeleton. For each fold, the bottom line voxels can be distinguished from the inner part of the fold surface because they do not split the skeleton background into two different local connected components [19] (Figure

1A). Then, the bottom line tag permits to define a topology-based augmentation, which conserves the bottom lines in all views but remove the inner part of some folds. The first view combines random $[-10,10]^\circ$ rotations over all axes and a 60% rolling cutout with only bottom lines kept inside the cutout volume. The second view combines random $[-10,10]^\circ$ rotations over all axes followed by a 60% rolling cutout with the whole skeleton conserved inside the cutout whereas only bottom values are kept outside the cutout volume. All views are then binarized. This topology-based augmentation forces the model to learn the sheet-based structure of the fold-based skeleton.

To decode SimCLR latent code, we freeze SimCLR weights and train a decoder whose input layer is the representation space. The decoder backbone is the one of the $\beta - VAE$ decoder, and the decoder loss is the cross entropy reconstruction error.

2.3 Identifying folding patterns

Characterizing folding shapes. To identify folding patterns, data are encoded to the latent space of both models and reduced to a 2-dimensions space with t-SNE algorithm. The reduction to two dimensions enables to get more hints of the learned representations and to analyze subjects groups more easily. A clustering is then performed with hierarchical affinity propagation (AP) algorithm [11]. One advantage of AP is that the number of expected clusters does not have to be precised. However it may output a very large number of clusters, making it difficult to understand from an anatomical point of view. Hence, following the method used in [20], we applied the algorithm in an iterative way until a maximum number of five clusters is found.

The analysis of the main anatomical characteristics of the clusters can be done either on the latent codes or on the input space based on cluster labels. The first method enables to understand the encoded characteristics in the latent dimensions. For $\beta - VAE$ we generated images corresponding to clusters' centroids from their latent codes which are next decoded. Then, we travelled between clusters through the latent space to analyse variations across dimensions. For SimCLR, we reconstructed latent representation of the nearest subject of each cluster centroid. The second method computes the local *per*-cluster averaging pattern in the input space [25].

3 Experiments and results

Datasets. We use HCP database³ in which MRI images were obtained with a Siemens Skyra Connectom scanner with isotropic resolution of 0.7 mm. We focused on the right hemisphere of 550 subjects. 80% of subjects were used for training and the remaining 20% were used for validation.

³ <https://www.humanconnectome.org/>

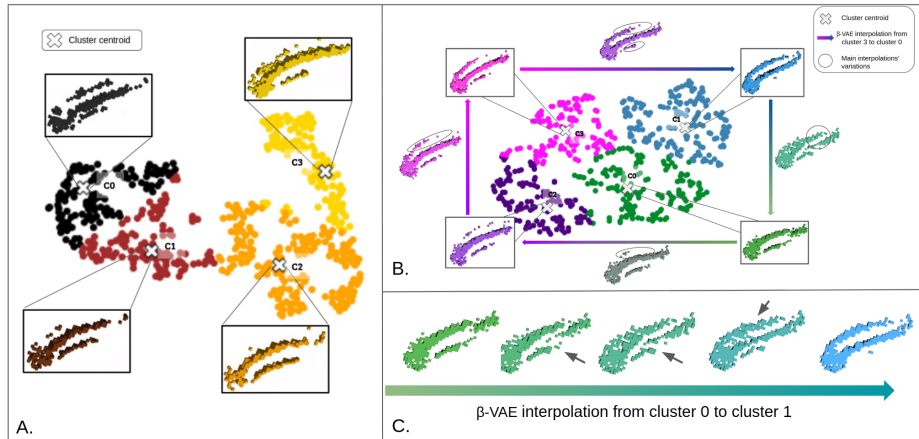


Fig. 2. β -VAE and SimCLR latent spaces analysis. (A.) t-SNE representation of SimCLR latent space. Insets are decoded latent codes of nearest neighbours for each cluster centroid. (B.) t-SNE representation of β -VAE latent space. Insets are decoded latent codes of cluster centroids and of interpolations between cluster centroids. Leftmost and rightmost patterns of (C.) are respectively the decoded latent code of cluster 0 and 1 centroids of the β -VAE model. Intermediate patterns of (C.) are obtained traveling through the latent space and then decoded.

Model implementation. Our β -VAE comprises fully convolutional encoder and decoder of symmetrical architectures with three convolutional blocks and two fully connected layers. The backbone of our SimCLR model is the DenseNet [14], followed by two fully connected projection overheads based on [9] benchmark on 3D MRI images. To adapt to our smaller input, we reduced the size of the DenseNet network down to two dense blocks. We call latent space, the representation space of the SimCLR model, which has a better representation quality than output space [6].

To find the best hyperparameters (size of the latent space for both models, β value for β -VAE and temperature τ for SimCLR), we performed a gridsearch where the best combination is chosen based on the loss value, the silhouette score on the latent space and the reconstruction abilities for β -VAE. We obtained $\beta=2$ (tested range 1-8) and $\tau=0.1$ (tested range 0.01-0.3), as well as a latent size of 4 (tested range 2-150) for both models, which enabled to balance between the model performance and the clustering quality. Training of 300 epochs lasted for approximately 1 hour and 2 hours for β -VAE and SimCLR respectively, on an Nvidia Quadro RTX5000 GPU.

Latent space structure. Fig. 2A and B present clustering results. The silhouette score with AP on the latent space is 0.17 and 0.42, respectively for β -VAE and SimCLR. It becomes 0.43 and 0.44 when applied to the t-SNE space, indicating a tendency towards a clustered distribution with close clusters. This

range of score is common when dealing with complex data such as neuroimaging modalities [16]. For both models, four clusters were identified but the organization of the latent space is different: β -VAE latent space seems to distinguish four groups of subjects, separated only with a thin boundary whereas SimCLR latent space is more structured and could be interpreted as a manifold, consistent with the biological reality of folding patterns.

Deciphering the patterns. Both models were able to produce reconstructions that are compliant with the inputs, presenting a simplified version of the scene which enables to focus and bring out the most important features (Fig. 2A and B). For SimCLR, the black cluster seems to have a small paracingulate, which is a pattern described in the literature [20]. The brown and orange cluster seems to correspond to two parallel sulci without a callosal sulcus. Based on the latent space organization, similar to a manifold, it is interesting to analyse the reconstructions in terms of evolution. The curvature of the longest sulcus becomes more bent from the black to the yellow cluster.

In the case of β -VAE, cluster 0 (green) shows another pattern defined in the literature [20]: a split anterior cingulate sulcus. Cluster 1 (blue) presents a long paracingulate, while the pink has a shorter paracingulate. Lastly, the indigo presents a slight paracingulate. In Fig. 2C, interpolating from one cluster to another shows that the latent space is continuous and regular, and we can progressively see the change of patterns as indicated by the arrows. More detailed and complete interpolations are presented in supplementary materials.

It is interesting to link these decoder outputs to the cluster average of the folding pattern based on the input space (figure 3). For SimCLR model, the black cluster average could correspond to a simple anterior cingulate. Subjects of the brown cluster could have a sketch of the paracingulate sulcus, which increases in length in the orange average to present two long parallel sulci. Finally, the yellow average also includes a sketch of a sulcus parallel to the anterior cingulate, but in the left part of the ROI, where it is not usually called a paracingulate sulcus in anatomical literature. Both methods, cluster averages and decoders, represent something different: in the first case, it is the geometrically-aligned average of all subjects in a cluster; in the latter case, it is the reconstruction of one representative subject from the latent space. They can converge either to the same (orange and yellow cluster) or to an apparently different (black cluster) representation.

For β -VAE, when comparing with centroids' generation, we find a similar shape for the green average (cluster 0): cingulate split in two. Conversely, for the blue cluster, based only on the average pattern, we could interpret a simple cingulate, but in the light of the reconstructions, the swollen anterior part could represent a paracingulate. Cingulate and paracingulate could be merged in the average representation due to positional variations among subjects. For indigo and pink folding averages, only the highly variable paracingulate pits are kept contrary to the decoder reconstructions. Thus, only the decoder part will permit

to get rid of the too complex inter-individual variability and to focus on the main shapes.



Fig. 3. *Representative patterns as cluster averages.* **A)** Description of typical folding structures in our ROI using the icbm152 average template. **B), C)** Local average sulci obtained for each cluster with β -VAE and SimCLR encodings respectively. Colors match cluster colors of Fig. 2.

4 Discussion and conclusion

Our work proposes several method contributions that can be useful for the community. We introduced topology-based augmentations in the SimCLR setting, which is directly applicable for studies working on skeletons or similar inputs [23]. It is all the more interesting as the augmentations used for contrastive views is still under investigation to understand what makes good views, especially in biomedical images. Moreover, we added a decoder to SimCLR and analyzed SimCLR and β -VAE reconstructions to recover folding patterns. Last, we proposed a preprocessing based on a mask, enabling to focus on the region of interest, while avoiding the disadvantages of parallelepipedic bounding boxes used in [12].

Our work also finds a structured latent space for the cingulate region with both models, β -VAE and SimCLR. The organization obtained with SimCLR seems more consistent with anatomical reality of folding patterns and can be linked to folding manifolds [18]. In return, the generative and regularization aspect of the β -VAE is a real lever to understand the learned representations and ease the analysis of this complex region.

To encourage a structured and well separated latent space, we wish to introduce in future cluster objectives in the learning phase both for generative models [7], and for contrastive models [4,5,17].

Another line of research will be to adapt our model further to the folding topology by developing other topology-based augmentations and by introducing other specific priors such as the geometry-based similarity measure between input samples [25].

Finally, we found both cluster averages and decoder outputs to be similar to known cingulate patterns that correlate with executive functions and psychiatric disorders. This similarity makes us firmly believe that such latent space structures could correlate with medically relevant parameters. This correlation will be the object of a follow-up study.

5 Acknowledgments

We thank colleagues from the Deep learning journal club for thorough discussions. We thank Zaccharie Ramzi for helping set up rotation augmentations. This research received funding from the European Union’s Horizon 2020 Research and Innovation Programme under Grant Agreement No. 945539 (HBP SGA3), from the FRM DIC20161236445, the ANR-19-CE45-0022-01 IFOPA-SUBA, the ANR-20-CHIA-0027-01 FOLDDICO. Data were provided in part by the Human Connectome Project funded by the NIH.

References

1. Borne, L., Rivière, D., Mancip, M., 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)
2. Borst, G., Cachia, A., Vidal, J., Simon, G., Pineau, A., Fischer, C., Poirel, N., Mangin, J.F., Houdé, O.: Folding of the anterior cingulate cortex partially explains inhibitory control during childhood: A longitudinal study. *Developmental Cognitive Neuroscience* **9**, 126–135 (Jul 2014)
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. Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep Clustering for Unsupervised Learning of Visual Features. In: *ECCV*. pp. 139–156 (2018)
5. Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. *arXiv:2006.09882* (Jan 2021)
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A Simple Framework for Contrastive Learning of Visual Representations. In: *Proceedings of the 37th International Conference on Machine Learning*. pp. 1597–1607. PMLR (Nov 2020)
7. Danks, D., Yau, C.: BasisDeVAE: Interpretable Simultaneous Dimensionality Reduction and Feature-Level Clustering with Derivative-Based Variational Autoencoders. pp. 2410–2420. PMLR (2021)
8. Duan, D., Xia, S., Rekik, I., Meng, Y., Wu, Z., Wang, L., Lin, W., Gilmore, J.H., Shen, D., Li, G.: Exploring folding patterns of infant cerebral cortex based on multi-view curvature features: Methods and applications. *NeuroImage* **185**, 575–592 (Jan 2019)
9. Dufumier, B., Gori, P., Battaglia, I., Victor, J., Grigis, A., Duchesnay, E.: Benchmarking CNN on 3D Anatomical Brain MRI: Architectures, Data Augmentation and Deep Ensemble Learning. *arXiv:2106.01132* (Jun 2021)

10. 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: MICCAI 2021 - Accepted Papers and Reviews (Sep 2021)
11. Frey, B.J., Dueck, D.: Clustering by Passing Messages Between Data Points. *Science* **315**(5814), 972–976 (Feb 2007)
12. Guillon, L., Cagna, B., Dufumier, B., Chavas, J., Rivière, D., Mangin, J.F.: Detection of Abnormal Folding Patterns with Unsupervised Deep Generative Models. In: *Machine Learning in Clinical Neuroimaging*. pp. 63–72 (2021)
13. Higgins, I., Matthey, L., Pal, A., Burgess, C.P., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. In: *ICLR* (2017)
14. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely Connected Convolutional Networks. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 2261–2269 (Jul 2017)
15. Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. *ICLR* (May 2014), <http://arxiv.org/abs/1312.6114>, arXiv: 1312.6114
16. Lebenberg, J., Mangin, J.F., Thirion, B., Poupon, C., Hertz-Pannier, L., Leroy, F., Adibpour, P., Dehaene-Lambertz, G., Dubois, J.: Mapping the asynchrony of cortical maturation in the infant brain: A MRI multi-parametric clustering approach. *NeuroImage* **185**, 641–653 (Jan 2019)
17. Li, Y., Hu, P., Liu, Z., Peng, D., Zhou, J.T., Peng, X.: Contrastive Clustering. *Proceedings of the AAAI Conference on Artificial Intelligence* **35**(10), 8547–8555 (May 2021), number: 10
18. Mangin, J.F., Lebenberg, J., Lefranc, S., Labra, N., Auzias, G., Labit, M., Guevara, M., Mohlberg, H., Roca, P., Guevara, P., Dubois, J., Leroy, F., Dehaene-Lambertz, G., Cachia, A., Dickscheid, T., Coulon, O., Poupon, C., Rivière, D., Amunts, K., Sun, Z.: Spatial normalization of brain images and beyond. *Medical Image Analysis* **33**, 127–133 (Oct 2016)
19. Mangin, J.F., Frouin, V., Bloch, I., Rigis, J., 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**(4), 297–318 (1995)
20. Meng, Y., Li, G., Wang, L., Lin, W., Gilmore, J.H., Shen, D.: Discovering cortical sulcal folding patterns in neonates using large-scale dataset. *Human Brain Mapping* **39**(9), 3625–3635 (Apr 2018)
21. Ono, M., Kubik, S., Abarnathey, C.D.: *Atlas of the Cerebral Sulci*. Thieme-Stratton Corp, Stuttgart : New York, 1er édition edn. (Jan 1990)
22. Provost, J.B.L., Bartrés-Faz, D., Paillère-Martinot, M.L., Artiges, E., Pappata, S., Recasens, C., Pérez-Gómez, M., Bernardo, M., Baeza, I., Bayle, F., Martinot, J.L.: Paracingulate sulcus morphology in men with early-onset schizophrenia. *British Journal of Psychiatry* **182**(3), 228–232 (Mar 2003)
23. Rao, H., Wang, S., Hu, X., Tan, M., Guo, Y., Cheng, J., Liu, X., Hu, B.: A Self-Supervised Gait Encoding Approach with Locality-Awareness for 3D Skeleton Based Person Re-Identification. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 1–1 (2021), arXiv: 2009.03671
24. Roy, A., McMillen, T., Beiler, D.L., Snyder, W., Patti, M., Troiani, V.: A pipeline to characterize local cortical folds by mapping them to human-interpretable shapes. Tech. rep. (Nov 2020), [bioarxiv:2020.11.25.388785](https://arxiv.org/abs/2020.11.25.388785)

25. 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)
26. Taleb, A., Loetzsch, W., Danz, N., Severin, J., Gaertner, T., Bergner, B., Lippert, C.: 3D Self-Supervised Methods for Medical Imaging (2006), arXiv: 2006.03829
27. Weiner, K.S., Golarai, G., Caspers, J., Chuapoco, M.R., Mohlberg, H., Zilles, K., Amunts, K., Grill-Spector, K.: The mid-fusiform sulcus: a landmark identifying both cytoarchitectonic and functional divisions of human ventral temporal cortex. *NeuroImage* **84**, 453–465 (Jan 2014)
28. White, L.E., Andrews, T.J., Hulette, C., Richards, A., Groelle, M., Paydarfar, J., Purves, D.: Structure of the human sensorimotor system. I: Morphology and cytoarchitecture of the central sulcus. *Cerebral Cortex* **7**(1), 18–30 (Feb 1997)
29. Yousry, T.A., Schmid, U.D., Alkadhi, H., Schmidt, D., Peraud, A., Buettner, A., Winkler, P.: Localization of the motor hand area to a knob on the precentral gyrus. A new landmark. *Brain* **120**(1), 141–157 (Jan 1997)
30. Yücel, M., Wood, S.J., Phillips, L.J., Stuart, G.W., Smith, D.J., Yung, A., Velakoulis, D., McGorry, P.D., Pantelis, C.: Morphology of the anterior cingulate cortex in young men at ultra-high risk of developing a psychotic illness. *The British Journal of Psychiatry* **182**(6), 518–524 (Jun 2003)