



HAL
open science

Detection of abnormal folding patterns with unsupervised deep generative models

Louise Guillon, Bastien Cagna, Benoit Dufumier, Joël Chavas, Denis Rivière,
Jean-François Mangin

► **To cite this version:**

Louise Guillon, Bastien Cagna, Benoit Dufumier, Joël Chavas, Denis Rivière, et al.. Detection of abnormal folding patterns with unsupervised deep generative models. MLCN workshop, held in conjunction of MICCAI, Sep 2021, Strasbourg, France. hal-03349112

HAL Id: hal-03349112

<https://hal.science/hal-03349112v1>

Submitted on 20 Sep 2021

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

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

Detection of abnormal folding patterns with unsupervised deep generative models

Louise Guillon¹[0000-0001-5716-0070], Bastien Cagna¹[0000-0002-3031-3935],
Benoit Dufumier^{1,2}[0000-0002-8253-2363], Joël Chavas¹[0000-0002-3559-5919],
Denis Rivière¹[0000-0002-1902-2213], and Jean-François
Mangin¹[0000-0002-1612-461X]

¹ Université Paris-Saclay, CEA, CNRS, NeuroSpin, Baobab, Gif-sur-Yvette, France
louise.guillon@cea.fr

² LTCI, Télécom Paris, IPParis, France

Abstract. Although the main structures of cortical folding are present in each human brain, the folding pattern is unique to each individual. Because of this large normal variability, the identification of abnormal patterns associated to developmental disorders is a complex open challenge. In this paper, we tackle this problem as an anomaly detection task and explore the potential of deep generative models using benchmarks made up of synthetic anomalies. To focus learning on the folding geometry, brain MRI are preprocessed first to deal only with a skeleton-based negative cast of the cortex. A variational auto-encoder is trained to get a representation of the regional variability of the folding pattern of the general population. Then several synthetic benchmark datasets of abnormalities are designed. The latent space expressivity is assessed through classification experiments between control’s and abnormal’s latent codes. Finally, the properties encoded in the latent space are analyzed through perturbation of specific latent dimensions and observation of the resulting modification of the reconstructed images. The results have shown that the latent representation is rich enough to distinguish subtle differences like asymmetries between the right and left hemispheres.

Keywords: Variational autoencoder · Brain architecture · Cortical folding · Anomaly benchmark · Anomaly detection

1 Introduction

The cortex folds *in utero* to form numerous furrows called sulci, which delimit circumvolutions. Cortical folding is related to cortical architecture (architectony and connectivity) [8,7] and can be impacted by developmental issues that lead to brain disorders [25,9]. The identification of folding patterns acting as markers of developmental brain diseases would be a major breakthrough facilitating early diagnosis. However, although cortical morphology embeds a topography of the sulci sufficiently consistent across subjects to enable the design of automatic recognition tools [5], the shapes of the sulci present a high diversity, which

hinders the modelling of the inter-individual variability necessary to define abnormalities [18]. Hence, the diversity of the folding pattern is often put aside and canceled out using spatial normalisation, namely warping all brains toward a template space.

Associations between folding patterns and developmental disorders have already been described. For instance, a very rare pattern called the Power Button Sign (PBS) has been linked to the epileptogenic zone of patients suffering from drug-resistant type 2 focal cortical dysplasia [17]. Similarly, it was demonstrated that the paracingulate sulcus morphology is correlated to hallucinations in patients suffering from schizophrenia [24]. Sulci shape deviations have also been observed in autism spectrum disorder (ASD) [2,11].

Once abnormal folding patterns linked to a pathology have been identified, automatic detection techniques can be developed using supervised learning. For instance, the PBS can be detected with a supervised classifier [4]. However, the upstream process of identifying and defining new patterns of interest is tedious and difficult as each individual has a unique cortical folding geometry and spotting a recurrent abnormality throughout a set of patients is very complex. An unsupervised tool designed to uncover cortical folding abnormalities and potential biomarkers would be an important lever to harvest the potential meaning of unusual folding patterns.

In this paper we propose a dedicated framework based on deep learning and we test its potential through the detection of synthetic unusual folding patterns. The automatic detection of abnormal folding patterns is a challenging task that has not yet been addressed in the field of neuroimaging. In this work, we relate this objective with the general field of anomaly detection. Anomaly and novelty detection aims at identifying samples that do not fit the normal data distribution [19]. A few years ago, anomaly detection methods evolved towards deep learning approaches and specifically unsupervised deep learning due to the ability to detect potentially unseen events. Auto-encoder (AE) based methods have been particularly studied as they infer a latent space of interest with much fewer dimensions than the input space, enforcing to learn only the most common features of the training data. There exists a broad range of AE-based models. An extensive review on the detection of epilepsy lesions in MRI patches can be found in [1]. Similarly, different methods dedicated to medical images have been compared in [3] leading to qualify the variational AE (VAE) architecture as the most efficient. Generative adversarial networks (GAN) have also been used in order to identify biomarkers in optical coherence tomography scans, reaching good performances [22,21]. Based on these initial results, a first framework was proposed for anomaly detection in computed tomography scans of 3D brains where anomalies consisted in labeled traumatic brain injuries [23]. β -VAE have also been successfully used to model the inter-individual variability in the mouse brain [14]. More recently, very promising self-supervised methods have been applied to anomaly detection problems in medical images [6], but these methods lack the generative aspect provided by GAN and VAE, which

is crucial in terms of explainability. All these works assessed images containing known lesions. Our aim however is to discover still unknown patterns linked to diseases, which leads to challenging evaluation issues. Therefore, this paper is focused on dedicated synthetic benchmark datasets.

In this paper, we propose the first technique aiming at bringing to light unusual and potentially abnormal folding patterns. For this purpose, we propose first a dedicated preprocessing leading to focus the learning on the cortical folding geometry of a specific region of interest (ROI). Then, like in [14], a β -VAE is trained on a set of control data sampling the general population to get a latent representation of the folding pattern distribution in this ROI. We also create several benchmark datasets simulating unusual regional folding patterns to assess the ability of our model to detect them. Finally, we analyse the latent space capacity to separate regional patterns from the two hemispheres.

2 Methods

2.1 Focusing on Folding Information

Brain MRIs contain diverse information that are not all relevant to study folding patterns. Our method therefore includes a crucial first step of data pre-processing based on the BrainVisa/Morphologist pipeline (<http://brainvisa.info>) [16]. This pipeline combines several steps such as bias correction, grey-white segmentation, and skeletonization to obtain a negative cast of the folding. Morphologist’s skeletons were used as input of our learning model. They were first defined in [15] and are obtained by skeletonization of the grey matter and cerebrospinal fluid union while preserving the topology. The result is 3D volumes with three values: inside of the brain, sulci skeletons and outside of the brain. The use of these simple images rather than raw MRIs puts the focus of learning on the folding geometry and discard a major confound related to the width of the sulci, which increases with local atrophies induced by aging or degenerative pathologies.

2.2 Generating Synthetic Brain Anomalies

One of our biggest challenges is the lack of consensual datasets of abnormalities to assess the approach. The examples mentioned in the introduction are either especially challenging in terms of shape (the PBS in epilepsy) or inter-subject local variability (Broca’s area or Superior Temporal Sulcus (STS) branches in autism), or correspond to a stratification of the population into several frequent patterns, which is not in the scope of anomaly detection (paracingulate sulcus). Therefore, in this paper, we focus on a 3D ROI of 23x37x36 voxels with 2mm isotropic resolution, localized in only one of these challenging areas, the STS branches. This ROI has been defined in each subject using affine normalisation to the classical MNI reference space. The localization of the ROI in the MNI space has been learned from the open access training dataset with annotated sulci of the Morphologist pipeline [5]. We have designed several dedicated

synthetic anomaly benchmark datasets in order to be able to evaluate the performances of our model (Fig.1).

Deletion: Our first benchmark dataset consists of skeletons in which we have randomly deleted one piece of sulcus, which is chosen among topologically elementary parts called simple surfaces and proposed by the Morphologist pipeline [15]. To be deleted, a simple surface must be completely within the ROI and made up of more than 1000 voxels, which corresponds to about 17% of the average number of skeleton voxels in the ROI. This arbitrary threshold aims at performing modification of the geometry beyond the normal anatomical variability observed in the population.

Random: A second benchmark dataset is composed of random ROIs of the same dimensions and overlapping the skeleton but localized in different positions in the cortex. This benchmark is expected to be very easily spotted as abnormal since the images are highly different. It is used to ensure that the model is able to identify inputs that are far away from the normal distribution and that what the model has learned is not only non-region specific features such as voxel proportion and sulci continuity.

Asymmetry: The last benchmark dataset corresponds to the same ROI but defined in the other hemisphere and flipped. The flip is defined from the inter-hemispheric plane of the MNI space after affine spatial normalisation. This benchmark has a biological interest as hemispheric asymmetry is still an intense field of research.




Fig. 1. *Generated anomaly benchmarks.* Benchmark Deletion: Original crop and its modified version. Benchmark Random: two examples of random crops. Benchmark Asymmetry: crop of right hemisphere and crop of left hemisphere flipped.

2.3 Learning a representation of the normal variability

An effective way to model population variability is through AE-based networks. These architectures learn to project input data onto a lower dimensional manifold, also called latent space and to reconstruct from this space the input image. Simple AE are known to have some drawbacks and particularly the lack of regularization of the latent space. To overcome this issue, VAE model was introduced [13], and later β -VAE was proposed [10]. Like classical AE, β -VAE

are composed of two parts: an encoder and a decoder but add a variational objective. Contrary to simple AE, an input from image space \mathcal{X} , is encoded as a distribution in a latent space \mathcal{Z} comprising L dimensions, leading to a twofold objective. First, the minimization of the reconstruction error of the input image. Second, the matching of the encoded distribution to a prior distribution, usually a Gaussian, which is done thanks to Kullback-Leibler divergence and enables to regularize the latent space. VAE is a $\beta - VAE$ with KL divergence weighted at 1. Thus, $\beta - VAE$ encoder θ and decoder ϕ are trained maximising the following objective:

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

where $p(\mathbf{z})$ is the prior distribution, a reduced centered Gaussian distribution in our work that is approximated with the posterior distribution $q_{\phi}(\mathbf{z}|\mathbf{x})$. Tuning β parameter enables to improve latent factors disentanglement [10].

Analysing the Latent Space. The analysis of the latent space to understand the meaning of the encoded features is capital to assess the potential of our model to highlight unusual folding patterns. As such, we first trained a $\beta - VAE$ on normal data only, for our model to learn to encode normal variability. Next, normal and benchmark data unseen during training are projected to the latent space to perform this analysis. The resulting latent codes are used to train a gradient boosting algorithm to classify normal versus synthetic abnormal samples. These three classifiers are used first to ensure that the latent representations are able to capture some relevant information regarding the folding patterns. Then, we can focus our analysis on the features contributing the most to the success of the classification using the generative power of the $\beta - VAE$, like in [14]. We travel throughout the latent space modifying only one of these features and observe the generated folding patterns.

3 Results

3.1 Datasets and Implementation

To learn the inter-individual variability of control subjects, the HCP database was used³. MRIs were obtained with a single Siemens Skyra Connectom scanner with a resolution of 0.7mmx0.7mmx0.7mm. In our work, we studied only the right hemisphere of 997 right-handed subjects with high quality result of the Morphologist pipeline. 547 subjects were used for the training and 150 for the validation of the $\beta - VAE$. The remaining 300 subjects were used to train classifiers, half of them being used to create synthetic abnormal patterns following each of the benchmark experiment. Two third of these 300 subjects were used to estimate the $\beta - VAE$ hyper-parameters using a grid-search driven by a classifier and the last third was used to explore the latent space organization.

The skeletons were spatially normalized with an affine transformation to the standard MNI space and were down-sampled to a voxel size of 2mm, which

³ <https://www.humanconnectome.org/>

is sufficient to preserve the folding geometry. Voxels were set to 0 for inside the brain (26% of voxels), 1 for outside the brain (65%) and 2 for the sulci skeleton (9%). As mentioned above, the input to the $\beta - VAE$ was a 3D ROI, whose bounding box in MNI space was learned from the BrainVISA open access training set of 64 annotated brains, in order to include the two posterior branches of the STS. This 3D ROI is made up by $23 \times 37 \times 36$ voxels extended to $40 \times 40 \times 40$ using 1-padding.

The complete pipeline is shown in figure 2. Our $\beta - VAE$ was composed of fully convolutional encoder and decoder with symmetrical architectures comprising three convolutional blocks and 2 fully connected layers. We did a gridsearch ($L=8-100$, $\beta=1-20$, ranges are based on previous works [14] and reconstruction abilities), where hyperparameters were chosen according to the classification performances of the Deletion classifier applied to 100 controls and 100 synthetic samples using a 5-fold stratified cross-validation. We selected $L=100$, $\beta=2$ and a learning rate of $2e-4$. Training lasted for 400 epochs on a Nvidia Quadro RTX 8000 GPU and was completed in roughly 1 hour.




Fig. 2. *Whole pipeline.* First, bounding boxes of sulci of interest are defined. The HCP database is processed with the Morphologist pipeline and cropped thanks to the bounding boxes. Crops are then downsampled before feeding the $\beta - VAE$.

3.2 Analysing Learned Folding Variability

First, we visually evaluated the reconstruction ability of our β -VAE. 2D slices of several reconstructed inputs are presented in figure 3. For the control set, deletion and asymmetry benchmarks, the reconstructions are approximate but retain the general geometry. The main sulci included in the ROI can be identified. In return, the random crops cannot be reconstructed by the decoder, which outputs an image looking like a disturbed configuration of STS branches. This suggests that the model has really learned the distribution of the specific geometry of this cortical region rather than a generic distribution covering any skeleton configuration.




Fig. 3. Reconstructions of test inputs. 2D sections from the 3D ROI presented in sagittal view at depth 18. First row: original images, second row: model outputs.

To evaluate our model latent space, three gradient boosting classifiers were trained using 50 controls and 50 synthetic samples from the test set, using a 5-fold stratified cross-validation. ROC curves are presented in figure 4.A). For each configuration, the AUC score is above chance. Highest performances are obtained on benchmark Random (AUC = 0.98), which was expected. This first result comforts the fact that our model is able to capture very obvious abnormalities. On benchmark Asymmetry, very good scores are also obtained (AUC = 0.85). Though for an inexperienced eye the difference between right hemisphere and flipped left hemisphere can seem subtle, this region is known to be asymmetrical in terms of length and tilt of the main sulci [20]. It confirms that our model is capable of representing specific anatomical structures included in the ROI. Finally, when deleting one large simple surface, results are slightly above chance (AUC=0.69) which indicates a potential of the latent space to detect such anomalies. Nevertheless further work is required using a larger test set to overcome potential limitations of this classifier experiment.

Using gradient boosting classifiers gives insights on the most decisive dimensions of the latent space, which depend on the benchmark. Figure 4.B) shows a visualization of the datasets using the two most important latent features or tSNE 2D manifold projection. For Random and Asymmetry benchmarks, two groups can be clearly identified. Surprisingly, in the tSNE visualization, the random crops are surrounded by controls, which is counter-intuitive relative to the Gaussian prior and does not fit with the plot performed from the two most discriminative dimensions. Further experiments with unbalanced datasets including only a small ratio of abnormality will help us clarify this point. The Suppression benchmark is clearly the most difficult classification experiment with the current latent space.

The last experiment consists in sampling latent vectors to travel along the most important dimensions according to the classifiers. Figure 4.C) shows the reconstruction provided by the decoder when following the dimension corresponding to the best discriminator of Asymmetry benchmark. All other dimensions were fixed at their mean value. As expected, the outputs look like approximated skeletons of the ROI. Subtle changes consistent during the travel can be observed: i.e. the upper part of the sulcus presented in pink on figure 4.C), called




Fig. 4. *Analyses of latent space.* **A)** ROC curves of GradientBoosting classifiers for classifications between controls and benchmarks. **B)** Distribution of data points using classifiers most important features and t-SNE. **C)** Travelling through the latent space.

the Sylvian fissure (SF), seems to shorten from the left to the right of the dimension. This observation is consistent with Asymmetry benchmark’s subjects distribution according to the two most important features on figure 4.B). Indeed, control subjects, i.e. right hemispheres correspond to higher values of the 88th dimension which is also suggested by the SF shortening on figure 4.C). This evolution is interesting as previous works demonstrated that the SF was shorter in the right hemisphere. Thus, this dimension could encode the length of the SF which is an important marker of laterality [12].

4 Discussion and Conclusion

In this paper we developed a framework that shall lead to the discovery of abnormal folding patterns beyond reach for human cognition because of the high inter-subject normal variability. Our main contribution is the design of synthetic benchmarks used to decipher the organization of the latent space used to model this variability. We have shown that the regional specificity of the folding pattern can be learned and used to detect some deviations from the norm. Our methods achieved to detect obvious and more subtle deviations (respectively with Random and Asymmetry benchmarks), but detection is harder for more complex ones such as Deletion benchmark. However, we don’t seek to detect benchmark subjects but rather to detect abnormal patterns. In the future we plan to perform further experiments to get more hints about the nature of the representation. Although some interesting features can be observed when varying latent vectors

dimensions' values, it is very difficult to apprehend all possible shifts looking at 2D sections because we deal with 3D images. Additionally, in some cases, we observe that the generated images are not realistic, preventing the interpretation. Generating synthetic anomaly benchmarks is a very useful initial step but induces biases that must be acknowledged. In our case, with the Deletion benchmark we made the hypothesis that anomalies could be linked to "missing" simple surfaces. This work constitutes a first step of proof of concept and enables to control the complexity of abnormalities, however an important limitation is the high dependency of the results on simulated data as our benchmarks are made of synthetic anomalies. Nevertheless we stress out that the model has learned only control data, thus is totally unsupervised in regards to anomaly data, and only the evaluation results depend on simulated data. In future works, the ultimate benchmarks will have to be built from large datasets of neurodevelopmental disorders and will aim at discovering actual abnormal folding patterns.

Acknowledgments

This project has received funding from the FRMDIC20161236445, the ANR-19-CE45-0022-01 IFOPASUBA, the ANR-14-CE30-0014-02 APEX the ANR-20-CHIA-0027-01 FOLDDICO. Data were provided in part by the Human Connectome Project funded by the NIH. This work was performed using HPC resources from GENCI-IDRIS (Grant 2020-AD011011929).

References

1. Alaverdyan, Z.: Unsupervised representation learning for anomaly detection on neuroimaging. Application to epilepsy lesion detection on brain MRI. Ph.D. thesis, Université de Lyon (Jan 2019)
2. Auzias, G., Viellard, M., Takerkart, S., Villeneuve, N., Poinso, F., Fonséca, D.D., Girard, N., Deruelle, C.: Atypical sulcal anatomy in young children with autism spectrum disorder. *NeuroImage: Clinical* **4**, 593–603 (Jan 2014). <https://doi.org/10.1016/j.nicl.2014.03.008>
3. Baur, C., Denner, S., Wiestler, B., Albarqouni, S., Navab, N.: Autoencoders for Unsupervised Anomaly Segmentation in Brain MR Images: A Comparative Study. arXiv:2004.03271 [cs, eess] (Apr 2020), arXiv: 2004.03271
4. Borne, L., Rivière, D., Cachia, A., Roca, P., Mellerio, C., Oppenheim, C., Mangin, J.F.: Automatic recognition of specific local cortical folding patterns. *NeuroImage* **238**, 118208 (Sep 2021). <https://doi.org/10.1016/j.neuroimage.2021.118208>
5. 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). <https://doi.org/10.1016/j.media.2020.101651>
6. Bozorgtabar, B., Mahapatra, D., Vray, G., Thiran, J.P.: SALAD: Self-supervised Aggregation Learning for Anomaly Detection on X-Rays. In: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M.A., Zhou, S.K., Racoceanu, D., Joskowicz, L. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. pp. 468–478. Lecture Notes in Computer Science, Springer

- International Publishing, Cham (2020). https://doi.org/10.1007/978-3-030-59710-8_46
7. Fernández, V., Llinares-Benadero, C., Borrell, V.: Cerebral cortex expansion and folding: what have we learned? *The EMBO Journal* **35**(10), 1021–1044 (May 2016). <https://doi.org/10.15252/emboj.201593701>
 8. Fischl, B., Rajendran, N., Busa, E., Augustinack, J., Hinds, O., Yeo, B.T., Mohlberg, H., Amunts, K., Zilles, K.: Cortical Folding Patterns and Predicting Cytoarchitecture. *Cerebral Cortex* **18**(8), 1973–1980 (Aug 2008). <https://doi.org/10.1093/cercor/bhm225>
 9. Guerrini, R., Dobyns, W.B., Barkovich, A.J.: Abnormal development of the human cerebral cortex: genetics, functional consequences and treatment options. *Trends in Neurosciences* **31**(3), 154–162 (Mar 2008). <https://doi.org/10.1016/j.tins.2007.12.004>
 10. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework (Nov 2016), <https://openreview.net/forum?id=Sy2fzU9gl>
 11. Hotier, S., Leroy, F., Boisgontier, J., Laidi, C., Mangin, J.F., Delorme, R., Bolognani, F., Czech, C., Bouquet, C., Toledano, E., Bouvard, M., Petit, J., Mishchenko, M., d’Albis, M.A., Gras, D., Gaman, A., Scheid, I., Leboyer, M., Zalla, T., Houe-nou, J.: Social cognition in autism is associated with the neurodevelopment of the posterior superior temporal sulcus. *Acta Psychiatrica Scandinavica* **136**(5), 517–525 (Nov 2017). <https://doi.org/10.1111/acps.12814>
 12. Idowu, O.E., Soyemi, S., Atobatele, K.: Morphometry, Asymmetry and Variations of the Sylvian Fissure and Sulci Bordering and Within the Pars Triangularis and Pars Operculum: An Autopsy Study. *Journal of Clinical and Diagnostic Research : JCDR* **8**(11), AC11–AC14 (Nov 2014). <https://doi.org/10.7860/JCDR/2014/9955.5130>
 13. Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. arXiv:1312.6114 [cs, stat] (May 2014), arXiv: 1312.6114
 14. Liu, R., Subakan, C., Balwani, A.H., Whitesell, J., Harris, J., Koyejo, S., Dyer, E.L.: A Generative Modeling Approach for Interpreting Population-Level Variability in Brain Structure. In: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M.A., Zhou, S.K., Racocceanu, D., Joskowicz, L. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. pp. 257–266. Springer International Publishing, Cham (2020)
 15. Mangin, J.F., Frouin, V., Bloch, I., Rgis, 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)
 16. Mangin, J.F., Rivière, D., Cachia, A., Duchesnay, E., Cointepas, Y., Papadopoulos Orfanos, D., Collins, L., Evans, A., Régis, J.: Object-Based Morphometry of the Cerebral Cortex. *IEEE transactions on medical imaging* **23**, 968–82 (Sep 2004). <https://doi.org/10.1109/TMI.2004.831204>
 17. Mellerio, C., Roca, P., Chassoux, F., Danière, F., Cachia, A., Lion, S., Naggara, O., Devaux, B., Meder, J.F., 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**(2), 500–507 (Sep 2014). <https://doi.org/10.1148/radiol.14140773>, publisher: Radiological Society of North America

18. Ono, M., Kubik, S., Abernathey, C.D.: Atlas of the cerebral sulci. G. Thieme Verlag ; Thieme Medical Publishers, Stuttgart; New York (1990), oCLC: 645306373
19. Pang, G., Shen, C., Cao, L., Hengel, A.v.d.: Deep Learning for Anomaly Detection: A Review. arXiv:2007.02500 [cs, stat] (Jul 2020), arXiv: 2007.02500
20. Rubens, A.B., Mahowald, M.W., Hutton, J.T.: Asymmetry of the lateral (sylvian) fissures in man. *Neurology* **26**(7), 620–620 (Jul 1976). <https://doi.org/10.1212/WNL.26.7.620>, publisher: Wolters Kluwer Health, Inc. on behalf of the American Academy of Neurology Section: Articles
21. Schlegl, T., Seeböck, P., Waldstein, S.M., Langs, G., Schmidt-Erfurth, U.: f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks. *Medical Image Analysis* **54**, 30–44 (May 2019). <https://doi.org/10.1016/j.media.2019.01.010>
22. Schlegl, T., Seeböck, P., Waldstein, S.M., Schmidt-Erfurth, U., Langs, G.: Unsupervised Anomaly Detection with Generative Adversarial Networks to Guide Marker Discovery. arXiv:1703.05921 [cs] (Mar 2017), arXiv: 1703.05921
23. Simarro Viana, J., de la Rosa, E., Vande Vyvere, T., Robben, D., Sima, D.M.: Unsupervised 3D Brain Anomaly Detection. In: BrainLes Workshop. Lima (Oct 2020)
24. The Australian Schizophrenia Research Bank, Garrison, J.R., Fernyhough, C., McCarthy-Jones, S., Haggard, M., Simons, J.S.: Paracingulate sulcus morphology is associated with hallucinations in the human brain. *Nature Communications* **6**(1), 8956 (Dec 2015). <https://doi.org/10.1038/ncomms9956>
25. Walsh, C.A.: Genetic Malformations of the Human Cerebral Cortex. *Neuron* **23**(1), 19–29 (May 1999). [https://doi.org/10.1016/S0896-6273\(00\)80749-7](https://doi.org/10.1016/S0896-6273(00)80749-7)