DETECTING RARE CORTICAL CONNECTIVITY AROUND THE HUMAN CENTRAL SULCUS: A DEEP LEARNING ANALYSIS OF 37,000+ TRACTOGRAPHIES

C Mendoza, N Vindas, J Chavas, J Laval, A Dufournet, P Guevara, V Frouin, D Rivière, J.-F Mangin

DETECTING RARE CORTICAL CONNECTIVITY AROUND THE HUMAN CENTRAL SULCUS: A DEEP LEARNING ANALYSIS OF 37,000+ TRACTOGRAPHIES

C. Mendoza[1], *N. Vindas*[1], *J. Chavas*[1], *J. Laval*[1], *A. Dufournet*[1], *P. Guevara*[2], *V. Frouin*[1], *D. Rivière*[1], *J.-F. Mangin*[1]

[1]Université Paris-Saclay, CEA, CNRS, BAOBAB, NeuroSpin, Gif-sur-Yvette, France
[2]Faculty of Engineering, Universidad de Concepción, Concepción, Chile

ABSTRACT

White matter (WM) and cortical folding are closely linked, exhibiting normal variability across individuals that reflects the diverse patterns of brain organization. However, identifying rare patterns remains a challenging task because they are largely hidden by the high inter-subject variability of cortical anatomy. In this work, we present the first deep learning framework for detecting uncommon arrangements of WM fiber tracts in a large population ($n = 37,140$). To this end, we employed a β-Variational AutoEncoder (β-VAE) to learn meaningful latent representations of local WM connectivity. Leveraging the latent space for a classification task on an annotated subset of subjects with an interrupted Central Sulcus (CS) achieved an area under the receiver operating characteristic curve above 0.9 across several β-VAE models. Furthermore, by analyzing the predicted probability of an interrupted CS, we identified a population-level transition from a continuous, stable CS to a more disorganized pattern characterized by branching folds and WM occupying the sulcal space. Code is available at github.com/neurospin-projects/bvae_fibers.

Index Terms— Tractography, Central Sulcus, β-VAE

1. INTRODUCTION

The human cerebral cortex has a complex folding pattern shaped by genetic, cellular, and biomechanical influences during development [1]. These folding patterns remain largely stable after birth and can serve as morphological proxies for developmental disorders. For instance, the Power Button Sign (PBS) [2], formed by a branch of the Central Sulcus (CS) and a segment of the Precentral Sulcus, has been proposed as a marker for type II focal cortical dysplasia.

Another interesting configuration of the CS is the presence of a bridging gyrus connecting the precentral and postcentral regions, a rare feature observed in approximately 1% of the population [3]. This pattern may represent an extreme case of the *pli de passage* [4]—a deep annectant gyrus generating a protrusion in the sulcal floor. Such disruptions of the principal fissures arise during the earliest stages of sulcal development and should therefore be considered fundamental and distinctive features of the human brain [4, 1]. It is also reasonable to expect corresponding reorganizations of white matter (WM) connectivity in regions showing folding disruptions, resulting in rare connectivity patterns that are not yet very well understood.

Recent advances in high-resolution magnetic resonance imaging (MRI) have enabled the *in-vivo* investigation of brain structure and function with unprecedented detail. However, despite significant methodological and algorithmic progress, identifying rare anatomical patterns remains challenging, as they are often obscured by the high inter-subject variability of cortical anatomy. Deep learning offers a promising framework for addressing this issue by modeling the normal variability within a population and identifying individuals who deviate from this normative pattern.

In this work, we introduce a deep learning–based framework that identifies rare connectivity patterns across a population by leveraging diffusion MRI (dMRI) tractography [5]. First, we computed density maps of local connectivity from fibers terminating around the right CS. Then, we trained a β-Variational AutoEncoder (β-VAE) to detect rare organizations of the WM through reconstruction error and classification based on the latent space. We validated our methodology on a dataset of subjects annotated with an interrupted right CS, serving as a proxy of unusual WM geometry. We demonstrate the potential of this approach by achieving a high area under the receiver operating characteristic curve (AUC) in classifying interrupted CS cases, even when these patterns are modeled as part of the normal variability within the population.

2. MATERIALS AND METHODS

2.1. Diffusion MRI tractography

We used dMRI data from 37,426 subjects in the UKB database [6]. The dMRI data were acquired with two b-values (1,000 and 2,000 *s/mm²*), each with 50 diffusion directions, and an isotropic voxel resolution of 2.0 *mm*. For each subject, probabilistic tractography based on second-order integration over Fiber Orientation Distributions (iFOD2) [7], multi-shell multi-tissue Constrained Spherical Deconvolution (CSD) [8], and Anatomically-Constrained Tractography (ACT) [9] was

calculated using MRtrix3 software [10]. The following parameters were used for fiber tracking: 10 million fibers, maximum angle of 45°, step size of 1 *mm*, minimum length of 10 *mm*, maximum length of 250 *mm*, and FOD amplitude cutoff of 0.06. Seed points were determined dynamically over the WM FOD image. Finally, SIFT2 [11] weights were computed for each tractography to provide a more biologically meaningful representation of structural connectivity.

2.2. Fiber bundle segmentation and density maps

Fiber bundle segmentation was performed by selecting fibers surrounding the right Central Sulcus (CS). For this purpose, the right CS was manually labeled in an independent dataset of 62 subjects [12], then the subjects were affinely transformed to the ICBM09c space to generate a binary mask of the right CS, which was dilated to account for inter-subject variability in the position of the CS, as was done for a regional study of folding pattern variability [13]. This mask was then affinely transformed from ICBM09c space to individual dMRI space. Subsequently, fibers with at least one endpoint inside the CS mask were selected (see Fig. 1-A). No fiber length threshold was applied, enabling the segmentation of fibers originating from the whole brain. Next, a density map of the fiber endpoints was computed using the *tckmap* command from MRtrix3, with each streamline contribution weighted according to its SIFT2 weight. The resulting density map was then multiplied by the CS region mask to retain only the densities around the right CS (see Fig. 1-B), and then affinely transformed to ICBM09c. Finally, density maps in ICBM09c were cropped using the bounding box of the CS mask, resulting in smaller maps (84, 68, 91) with a voxel resolution of 1 *mm*. Quality control (QC) was performed by excluding subjects with data quality issues identified during the diffusion MRI processing pipeline (see section 2.1). Additionally, the number of segmented fibers was used to further exclude subjects with segmentation issues. As a result, a total of 37,140 subjects were included in the subsequent analysis.

2.3. Semi-automatic annotation of interrupted CS

UKB subjects exhibiting interrupted CS were identified using a human-guided active learning. From a representation space that captured the geometric characteristics of sulcal shapes within the CS region [14], an initial logistic model was trained using five interrupted versus five continuous CS. The model's predictions of interruption were then applied to other subjects, and the twenty individuals with the highest predicted probability of interruption were selected for manual review. A human expert updated the annotated dataset when subjects truly exhibited interrupted sulci. Some of the remaining subjects were assigned to the continuous group. After 140 iterations, the process gradually enhanced the model's accuracy and expanded the labeled database, resulting in the identification of 183 interrupted subjects with available dMRI data.

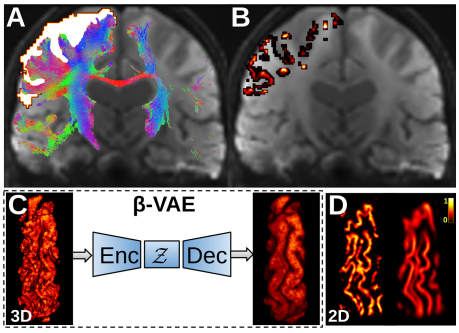

Fig. 1. (A) Segmentation of fibers terminating on the right CS mask. (B) Density map of the fiber bundle endpoints. (C) Simplified illustration of the β-VAE for one validation subject ($D = 512$, $\beta = 4$). (D) 2D cuts of the figures in (C) for the input (left) and reconstruction (right). Abbreviations: Enc: Encoder, $\mathcal{Z}$: Latent space, Dec: Decoder.

2.4. β-VAE architecture

The β-VAE is a stochastic encoder/decoder framework. The encoder maps the input to a distribution in a latent space of D dimensions. Then, a vector is sampled from this distribution and reconstructed by a decoder. The optimization process minimizes the reconstruction error between the input and the output, as well as the Kullback-Leibler Divergence (KLD) between the latent space distribution and a standard normal distribution. The parameter β weights the KLD, with higher values of β leading to a more disentangled latent space at the cost of poorer reconstructions.

In this work, we used the β-VAE implementation from [13], which consists of 3D convolutions of depth 3. Mean Squared Error (MSE) loss was used to process images with continuous values, learning rate and batch size were respectively fixed to 0.0002 and 16, and we determined the number of epochs based on early stopping as done in [13]. The input to the β-VAE are the cropped density images, which are log-transformed and min-max scaled before training to ensure proper reconstructions. Once a model is trained, rare patterns of cortical connectivity can be assessed through high reconstruction error or classification based on the latent space.

2.5. β-VAE for rare cortical connectivity detection

We evaluated the detection power of the β-VAE using the annotated subset of subjects with interrupted CS (see section 2.3). Specifically, we assessed whether this rare pattern can be detected, even when it is modeled as part of the normal variability of the population. For this purpose, we leveraged the reconstruction error and latent space of the β-VAE. The data was split according to Table 1, where *Validation 1* was used to compare the reconstruction error distribution of interrupted CS to the rest of the set and *Validation 2* was used to assess classification performance from the latent space us-

ing a Support Vector Machine (SVM) with a linear kernel and regularization parameter $C = 0.01$ as done in [14]. *Validation 2* also contained 92 continuous CS, which were verified by a human expert. The β-VAE was trained over a grid of parameters for β and the latent dimension *D* (refer to Table 2 for details). Some unknown interrupted CS may be hidden within Train and *Validation 1*. Mann-Whitney U test [15] was applied to assess significant differences between reconstruction errors in *Validation 1* across the various β-VAE models. Bonferroni correction [16] was applied to account for multiple comparisons using $\alpha = 5\%$ (36 comparisons).

Table 1. Number of subjects for each dataset. Also, the number of subjects with interrupted CS within each set is shown.

	Train	Validation 1	Validation 2
n	29,566	7,391	183
Interrupted CS in the set	74	18	91

2.6. Exploring the latent space organization

We selected an appropriate β-VAE model based on the three sets shown in Table 1 and a sum of criteria: (1) the average reconstruction error on *Validation 1*, normalized between 0 and 1, (2) the distance of reconstruction error distributions between the Train set and *Validation 1* to minimize the overfitting (3) the AUC of the mean receiver operating characteristic curve (ROC) from the *Validation 2* classification task and (4) the distance in reconstruction error distributions between interrupted and continuous CS subjects in *Validation 2*. To measure the distance between distributions, we used the statistic from the two-sample Kolmogorov-Smirnov (KS) test [17]. Criteria (1) and (2) were inverted to align their directionality with (3) and (4), enabling all metrics to be jointly maximized. Then, using the latent space of the selected β-VAE model, an SVM was trained on *Validation 2* to estimate the probability of interrupted CS for all subjects, and a four-dimensional UMAP was fit to reduce the dimensionality of the latent space for visualization. The SVM probability estimation was used to color the UMAP projection. Finally, the main variability traits of the population were explored by averaging data within probability bins representing 20% intervals.

3. RESULTS

3.1. Results for reconstruction error in Validation 1

When examining the reconstruction error distribution in *Validation 1*, we did not observe substantial differences, which was expected since the model had seen instances of interrupted CS during training. However, for a few β-VAE parameter configurations, the MSE was significantly higher for the interrupted CS group compared to the rest of the subjects ($p < 0.001$, Mann-Whitney U test). Fig. 2 shows the reconstruction error distribution corresponding to the largest significant difference between interrupted CS and the rest of the subjects of *Validation 1*.

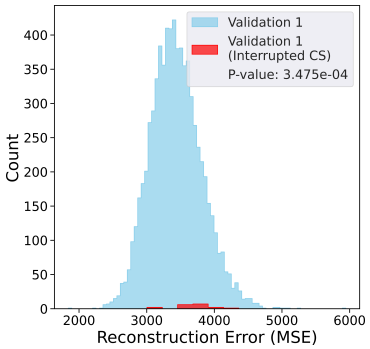

Fig. 2. Reconstruction error in *Validation 1* ($\beta = 4$, $D = 64$).

3.2. Results for classification task in Validation 2

We evaluated the performance of classifying interrupted CS in *Validation 2* using the latent space representations from different β-VAE models (see Table 2). The results showed an AUC above 0.8 for most combinations of β and latent dimension D, reaching up to 0.93 for some parameter settings. We observed that a $\beta = 32$, resulted in poor classification performance, likely due to the loss of relevant information, also producing blurrier reconstructions.

Table 2. Results for the classification in *Validation 2*. Linear SVM with Stratified K-Fold Cross Validation (K=5) was used. AUC of the mean ROC curve is shown.

β \ D	**64**	**128**	**256**	**512**	**1024**	**2048**
1	0.92	0.91	0.88	0.88	0.90	0.92
2	0.93	0.93	0.88	0.88	0.92	0.91
4	0.93	0.89	0.92	0.92	0.92	0.93
8	0.91	0.93	0.93	0.93	0.90	0.86
16	0.88	0.88	0.85	0.80	0.75	0.73
32	0.39	0.34	0.37	0.38	0.38	0.39

3.3. UMAP projection of the latent space

Based on the criteria described in section 2.6, we selected the β-VAE configuration with $\beta = 4$ and $D = 512$ as an appropriate model (Fig. 1 shows the reconstruction of a subject from the *Validation 1*). Then, based on a linear SVM trained on the latent space from *Validation 2*, we computed the probability of having an interrupted CS for all the subjects not in our annotated dataset. We found an interesting distribution of the data when plotting the first and fourth UMAP dimensions (see Fig. 3). In this figure, each subject is color-coded according to its probability, ranging from black (low probability) to

yellow (high probability), and revealing two hot spots: one in the leftmost region and one in the rightmost region of the UMAP distribution. All subjects annotated as interrupted CS are shown in red, and it can be seen that they were more biased toward the hot spots of higher probability. This tendency may suggest that the latent space encoded by the β-VAE captures features associated with the interruption of the CS, and the presence of two distinct hot spots may indicate a continuum of similarity captured by the model.

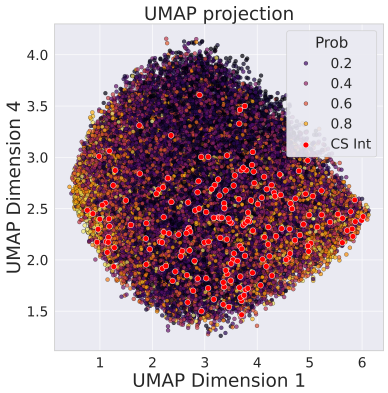

Fig. 3. UMAP of the latent space ($\beta = 4$, $D = 512$).

3.4. Detection of rare connectivity patterns by probability

We grouped the subjects not annotated as interrupted CS into 20% probability bins and computed the corresponding average density maps. These maps were smoothed, thresholded at various levels, and meshed to facilitate visualization. The resulting averages are shown in orange in Fig. 4-A, while the group average CS is overlaid in purple to add interpretability. A transition can be observed from a continuous CS to a more disorganized pattern, where the WM connectivity begins to invade the sulcal space, along with the merging of the CS with adjacent folding branches. Note the transition from a high number of subjects in the first probability bin ($n = 14,989$) to a lower number in the last bin ($n = 2,112$), which is consistent with our aim of finding rare WM connectivity patterns in the whole population. After examining the top 1,200 subjects with the highest probability of having an interrupted CS, we identified nine potentially new subjects not included in our annotated dataset (see Fig. 4-B, the first two subjects from left to right). Also, several subjects not fully interrupted but with a prominent *pli de passage* were observed as shown in the most right subject from Fig. 4-B. Finally, most of the subjects examined presented rare configurations in which the CS merged with small folding branches (see Fig. 4).

4. DISCUSSION AND CONCLUSIONS

In this work, we propose the first framework for detecting rare cortical connectivity patterns across a large population of

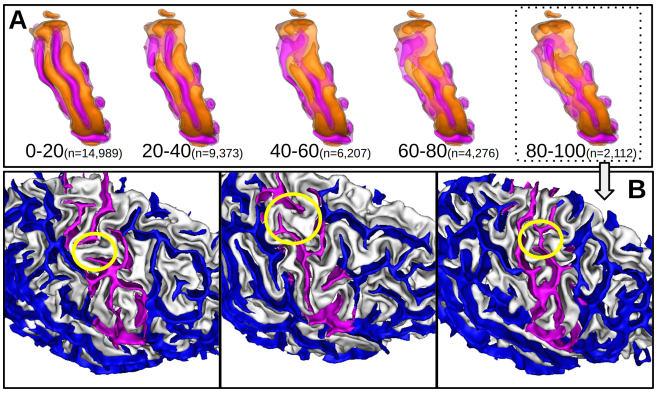

Fig. 4. (A) Population averages of the density maps (orange) and folds (purple) for different probability bins; the number of subjects in each bin is shown in parentheses. (B) New subjects with interrupted CS and prominent *pli de passage* were found in the group with the highest probability (80%-100%)

subjects. Our methodology is based on a β-VAE, designed to learn meaningful representations of WM connectivity through an encoder/decoder architecture. To handle the high amount of information embedded in tractograms, we computed a density map from the fiber bundle endpoints, making the input data suitable for standard 3D convolutional neural networks.

Our main results are twofold. First, we showed that we can detect rare connectivity patterns even when they are modeled as part of the normal population. This was demonstrated by the high AUC score obtained in *Validation 2* for a wide range of β-VAE parameters. Reconstruction error was less effective at detecting this rare pattern, possibly because the model had seen some examples during training. However, for some β-VAE models, we observed a significantly higher reconstruction error in the interrupted CS group.

Second, we demonstrated that the model captures meaningful latent representations of WM connectivity, as evidenced by the reconstruction of a single subject in Fig. 1, as well as the high AUC achieved in the classification of interrupted CS. Furthermore, by using the probability bins derived from the SVM, we observed a population-level transition from a continuous, stable CS to a rarer, more disorganized pattern. When examining the group with the highest probability of having an interrupted CS, we identified nine new potentially interrupted subjects and many exhibiting a prominent *pli de passage*, which could correspond to the thin fold line observed in the group average. The SVM may also be influenced by fold branching, which could locally disrupt WM connectivity. In future work, we will incorporate a fiber-length threshold to focus on distinct profiles of local connectivity. Due to the close anatomical relationship between superficial WM and cortical folding, it would be interesting to generate density maps restricted to short association fibers and evaluate whether these provide a more informative input.

5. ACKNOWLEDGMENTS

This research was conducted using the UK Biobank resource under application number 64984. This project has been funded by ANR via the Audace research program of the CEA, the IHU ICE (ANR-23-IAHU-0010), the PEPR Digital Health via BHT (ANR-22-PESN-0012) and StratifyAging (ANR-22-PESN-0010), and RHU-PsyCARE (ANR-18-RHUS-0014), by AXA foundation for the project premaIA, by the "Fondation de France" for the project TSA-ideogrammes. This project was provided with computing HPC and storage resources by GENCI at TGCC thanks to the grant 2024-A0170313800 on the supercomputer Joliot Curie's ROME partition. This project was provided with computer and storage resources by GENCI at IDRIS thanks to the grant AD010316018 on the supercomputer Jean Zay's V100 partition. This work was funded by the National Agency for Research and Development (ANID) / Scholarship Program / DOCTORADO BECAS CHILE / 2024 - 72240205.

6. REFERENCES

[1] H. de Vareilles, D. Rivière, J.-F. Mangin, and J. Dubois, "Development of cortical folds in the human brain: An attempt to review biological hypotheses, early neuroimaging investigations and functional correlates," *Developmental Cognitive Neuroscience*, vol. 61, pp. 101249, June 2023.

[2] C. Mellerio, P. Roca, F. Chassoux, F. Danière, A. Cachia, S. Lion, O. Naggara, B. Devaux, J.-F. Meder, and C. Oppenheim, "The power button sign: A newly described central sulcal pattern on surface rendering mr images of type 2 focal cortical dysplasia," *Radiology*, vol. 274, no. 2, pp. 500–507, Feb. 2015.

[3] J.-F. Mangin, Y. Le Guen, N. Labra, A. Grigis, et al., ""Plis de passage" deserve a role in models of the cortical folding process," *Brain Topography*, vol. 32, no. 6, pp. 1035–1048, Oct. 2019.

[4] D. J. Cunningham, "An address delivered at the opening of the section of anatomy and physiology," *BMJ*, vol. 2, no. 1544, pp. 277–283, Aug. 1890.

[5] P. J. Basser et al., "In vivo fiber tractography using DT-MRI data," *Magnetic Resonance in Medicine*, vol. 44, no. 4, pp. 625–632, 2000.

[6] C. Sudlow, J. Gallacher, N. Allen, et al., "UK biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age," *PLOS Medicine*, vol. 12, no. 3, pp. e1001779, Mar. 2015.

[7] J.-D. Tournier, F. Calamante, and A. Connelly, "Improved probabilistic streamlines tractography by 2nd order integration over fibre orientation distributions," *Proc. Intl. Soc. Mag. Reson. Med. (ISMRM)*, vol. 18, 01 2010.

[8] B. Jeurissen, J.-D. Tournier, T. Dhollander, et al., "Multi-tissue constrained spherical deconvolution for improved analysis of multi-shell diffusion MRI data," *NeuroImage*, vol. 103, pp. 411–426, Dec. 2014.

[9] R. E. Smith, J.-D. Tournier, F. Calamante, and A. Connelly, "Anatomically-constrained tractography: Improved diffusion MRI streamlines tractography through effective use of anatomical information," *NeuroImage*, vol. 62, no. 3, pp. 1924–1938, Sept. 2012.

[10] J.-D. Tournier, R. Smith, D. Raffelt, R. Tabbara, T. Dhollander, M. Pietsch, D. Christiaens, B. Jeurissen, C.-H. Yeh, and A. Connelly, "MRtrix3: A fast, flexible and open software framework for medical image processing and visualisation," *NeuroImage*, vol. 202, pp. 116137, Nov. 2019.

[11] R. E. Smith, J.-D. Tournier, et al., "SIFT2: Enabling dense quantitative assessment of brain white matter connectivity using streamlines tractography," *NeuroImage*, vol. 119, pp. 338–351, Oct. 2015.

[12] L. Borne, D. Rivière, M. Mancip, and J.-F. Mangin, "Automatic labeling of cortical sulci using patch- or CNN-based segmentation techniques combined with bottom-up geometric constraints," *Medical Image Analysis*, vol. 62, pp. 101651, May 2020.

[13] L. Guillon, J. Chavas, A. Bénézit, M.-L. Moutard, P. Roca, C. Mellerio, C. Oppenheim, D. Rivière, and J.-F. Mangin, "Identification of rare cortical folding patterns using unsupervised deep learning," *Imaging Neuroscience*, vol. 2, Feb. 2024.

[14] J. Laval, J. Chavas, V. Troiani, W. Snyder, M. Patti, M. Moyal, M. Plaze, A. Cachia, Z. Yi Sun, V. Frouin, P. Gori, D. Rivière, and J.-F. Mangin, *Towards a Foundation Model for Cortical Folding*, pp. 78–88, Springer Nature Switzerland, Dec. 2024.

[15] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *The Annals of Mathematical Statistics*, vol. 18, no. 1, pp. 50–60, Mar. 1947.

[16] W. Haynes, *Bonferroni Correction*, pp. 154–154, Springer New York, 2013.

[17] J. L. Hodges, "The significance probability of the smirnov two-sample test," *Arkiv för Matematik*, vol. 3, no. 5, pp. 469–486, Jan. 1958.