

Geometric-to-Semantic Spherical Transfer Learning for Cortical Sulci Labeling

Saeb Tounsi¹, Joël Chavas¹, Pietro Gori², Vincent Frouin¹, Denis Rivière¹, and Jean-François Mangin¹

¹ Paris-Saclay University, CEA, NeuroSpin, Baobab, Saclay, France

² LTCI, Télécom Paris, Institut Polytechnique de Paris, France
saeb.tounsi@gmail.com

Abstract. Deep learning on cortical surfaces faces a dilemma: capturing the complex topology of over 60 nomenclature-dependent sulci per hemisphere requires high-capacity models, yet the extreme scarcity of expert annotations ($N = 62$ subjects) inevitably causes overfitting. Standard supervised approaches fail to generalize on this data-scarce regime, particularly for variable and small sulci where topological ambiguity is high. To overcome this limitation, we introduce a Geometric-to-Semantic Spherical Transfer Learning framework. First, we leverage massive unlabeled data (UK Biobank with $\approx 30,000$ subjects) to pre-train a spherical encoder using a locally-optimized strategy. By relying solely on continuous surface features (curvature and depth), the relevance of this pre-training is confirmed by the model’s ability to detect localized and rare topological traits, such as sulcal interruptions. The downstream labeling task, however, introduces extracted sulcal fundi (lines) as an explicit semantic input. To bridge this dimensional domain gap (from purely geometric to semantic) without causing catastrophic forgetting, these anatomical lines are integrated into the pre-trained backbone via a soft-initialized Topological Prior Injector. Our experiments demonstrate that this approach outperforms fully supervised baselines trained from scratch, achieving a mean Dice of 0.77. Crucially, a local analysis reveals that the self-supervised geometric priors yield the largest performance gains on variable and tertiary sulci (up to 14.8%), confirming that learning the cortex shape is highly beneficial for identifying its rarest parts.³

Keywords: Transfer learning · Sulci labeling · Cortical Sulci · Self-supervised learning · Cortical surface

1 Introduction

The human cerebral cortex is characterized by a complex pattern of convolutions, where sulci serve as essential landmarks for structural and functional brain mapping. Accurate identification of these folds is crucial for understanding neurodevelopmental trajectories and improving the localization of functional areas.

³ https://github.com/neurospin-projects/2026_stounsi_slabin2

However, cortical sulci exhibit extreme inter-individual variability in terms of morphology, spatial location, length, and depth.

Historically, sulcal labeling has served two primary objectives: driving inter-subject brain alignment and enabling detailed morphometric studies. While the former can tolerate some geometric approximation, the latter demands high topological precision. The automated recognition of cortical sulci relied on volumetric approaches. Pioneering frameworks, such as BrainVISA [17], used a skeleton representation of folds and graph-based models to directly extract the topology of cortical folds from voxel space [1]. While these methods established the foundational principles of sulcal identification, methods working in 3D volumetric space often struggle to capture the intrinsic topology of the cortex.

Consequently, the field has transitioned toward surface-based representations (triangular meshes), popularized by pipelines like FreeSurfer [7]. Modeling the cortex as a 2-manifold is now considered essential for accurately characterizing morphology and aligning anatomy across subjects [18,5,7]. This shift allows for the analysis of features, such as curvature and sulcal depth, directly on the native geometry of the cortex, a prerequisite for precise topographical labeling [8].

The application of Geometric Deep Learning, particularly Spherical CNNs [3] and Graph Convolutional Networks (GCNs) [13], has shown promise in parcelating the cortical surface. However, exhaustive sulcal labeling remains an open challenge due to two critical constraints. First, state-of-the-art spherical U-Nets and similar surface-based architectures typically restrict their scope to a small subset of sulci (or only primary ones) [8,11,14]. For instance, recent work by Hao et al. [8] focuses on 6 to 13 sulci within specific regions (like the lateral prefrontal cortex) rather than a dense, whole-brain labeling of 63 structures as in [1]. Second, these models are data-hungry: robust performance typically requires massive cohorts (e.g., $N > 1400$ in [14]), making them inapplicable in expert-driven regimes where the complete annotated data is scarce ($N = 62$). Furthermore, the recent paradigm shift toward foundation models [19,9,4] in neuroimaging have demonstrated the power of learning universal representations from massive datasets, these models currently lack native support for surface-based topologies, which limits their direct application to highly convoluted cortical surfaces.

To bridge this gap, we propose a Spherical Transfer Learning framework that enables exhaustive identification of a comprehensive set of surface-consistent sulci in a severely data-constrained regime ($N = 62$). Our approach begins with a locally-optimized self-supervised pre-training on **30,000** subjects of the UK Biobank [20]. By forcing the network to focus on local surface patterns through targeted augmentations, it acquires “geometric intuition” sensitive to fine-grained anatomical nuances, such as sulcal interruptions. To transition this knowledge to the downstream task, we introduce a soft-initialized Topological Prior Injector (TPI) that bridges the domain gap between the two-channel geometric pre-training and the three-channel semantic labeling input. TPI allows the model to progressively integrate semantic information without suffering from catastrophic forgetting. This framework is supported by a surface-native pipeline that faithfully projects expert volumetric knowledge (as defined in [1] and [15]) onto the




Fig. 1. Overview of the proposed Spherical Transfer Learning framework. **(A)** Volumetric-to-Surface data preparation pipeline extracting 1D sulcal lines and projecting labels via Geodesic Voronoi. **(B)** Hybrid Spherical U-Net architecture. TPI projects the 3-channel input (C, S, L) into the pre-trained 2-channel encoder.

spherical mesh, providing a unified workflow for the characterization of a comprehensive set of anatomically consistent sulci.

2 Proposed Methodology

2.1 Data Preparation and Surface-Based Pipeline

Dataset and Fundi Extraction: The dataset used in this study comprises the right hemispheres of 62 subjects with expertly annotated cortical sulci, publicly available at https://brainvisa.info/data/sulci_database/base_62/2019/. We developed a strictly surface-native pipeline extracting sulcal lines from white matter meshes via TRACE [12]. Parameters were rigorously calibrated to exclusively isolate deep sulcal fundi, rejecting superficial dimples and wall ascensions.

Volumetric-to-Surface Label Projection: The ground-truth volumetric labels are subsequently mapped onto these 1D surface lines via a projection process based on geodesic Voronoi diagrams using a Fast Marching algorithm. To prevent aberrant semantic overlaps, such as a label improperly crossing a gyrus, this propagation is strictly constrained by the white matter mask and bounded by a maximum empirical geodesic distance threshold of 3.5 mm.

Spherical Input and Topological Regularization: To provide a standardized input structure for the neural network, the native meshes from FreeSurfer [7] are parameterized onto a unit sphere and resampled onto a regular icosahedron of order 6 ($\approx 42,000$ vertices). We use the unaligned spherical parameterization to preserve local intrinsic folds, deliberately avoiding standard non-linear registration that could distort sulcal geometry. This choice increases the labeling

difficulty, yet it is essential because folding patterns are not all compatible for a diffeomorphism-based alignment.

To prevent the emergence of artificial gaps within the sulcal fundi or the loss of local structural information after interpolation, we implemented a conservative projection strategy akin to graph-based spatial max-pooling: each target icosahedral vertex aggregates the native labels within its receptive field and retains the non-zero majority value. This transfer is immediately followed by graph-based mathematical morphology operations (specifically, a spatial closing via majority voting coupled with local erosion) applied iteratively in both the native and icosahedral spaces. The anatomical viability was validated through expert quality control on the selected 50 surface-consistent sulci.

2.2 Locally-Optimized Spherical Pre-training (SSL)

We employ a VGG-style spherical convolutional encoder to align with standard state-of-the-art U-Net architectures [23] used for cortical surface segmentation, the primary focus of this work being the adaptations made to maximize representation transferability. The network cascades spherical convolution blocks with local pooling operations. Unlike standard SSL frameworks that rely on Global Average Pooling (GAP) to foster spatial invariance, we directly flatten the final convolutional output into a high-dimensional vector. Because GAP suboptimally suppresses localization essential for labeling, omitting it ensures a more exhaustive spatial representation. The encoder is pre-trained on 30000 subjects of UK Biobank using the Barlow Twins self-supervised objective[22]. We generate augmented views using random spherical rotations with two structural occlusion strategies: random vertex dropout and contiguous patch masking. Although a demographic gap exists between the pre-training cohort (40–69 years) and the downstream dataset (predominantly 20–35 years), sulcal topology remains highly stable throughout adulthood, allowing the geometric priors to transfer effectively.

To verify that this globally trained latent space successfully encodes localized geometrical variations, we selected models by evaluating the frozen representations via linear probing on standard tasks (age regression, sex classification) and two topological ones:

- Central Sulcus (CS) Interruption: A rare, highly localized anatomical anomaly ($< 1\%$). We evaluated the model on 207 interrupted cases with 207 control cases identified using ChampollionV0 [10] classifier and verified by experts.
- Paracingulate Sulcus (PCS) Presence: A more balanced trait located in the anterior cingulate cortex [2,21]. Labels were generated following the procedure in [6]. To ensure label quality, we strictly filtered the dataset, retaining only high-confidence predictions (> 0.8 for presence, < 0.2 for absence).

2.3 Downstream Architecture: Transfer to the Hybrid U-Net

The identification task requires a 3-channel input as shown in Fig. 1 (curvature C , sulcal depth S , and extracted sulcal lines L as topological priors), which creates a dimensional mismatch with the 2-channel pre-trained SSL encoder (which

only uses C and S to avoid the trivial solutions that explicit sulcal lines would provide in the topological tasks). Naively expanding the first convolutional layer would corrupt the pre-learned filters and cause catastrophic forgetting.

To bridge this gap, we insert a **Topological Prior Injector (TPI)** directly before the encoder, implemented as a 1×1 convolution. Let $X \in \mathbb{R}^{V \times 3}$ be the input matrix for V vertices. It projects X into the 2-channel space $Y \in \mathbb{R}^{V \times 2}$ expected by the backbone:

$$Y = XW + b \quad \text{and} \quad W_{init} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ \epsilon_1 & \epsilon_2 \end{pmatrix}$$

where $W \in \mathbb{R}^{3 \times 2}$ and $b \in \mathbb{R}^2$.

To protect the pre-trained weights from abrupt distribution shifts, the novelty lies in the soft initialization of W with $\epsilon_1, \epsilon_2 \ll 1$. Initially, the network processes the exact geometric distribution it learned during pre-training ($Y_1 \approx C$, $Y_2 \approx S$). As fine-tuning progresses, the TPI smoothly integrates the semantic lines without abruptly shifting the input distribution.

We construct the Hybrid U-Net by grafting a randomly initialized, symmetric spherical decoder onto the pre-trained backbone. To handle small batch sizes, we used Group Normalization throughout the network. During fine-tuning, we apply differential learning rates: a low learning rate for the encoder to preserve its learned geometric manifold, and a standard rate for the TPI and decoder. Finally, spatial skip connections are added : they directly route the finely resolved local geometric features captured by the early SSL layers to the decoder, enforcing precise, dense labeling boundaries.

3 Experimental Setup and Evaluation Strategy

3.1 Baselines and Comparative Models

Dataset and Evaluation Strategy: We evaluate our framework on a cohort of 62 subjects with surface-native sulcal labels (as detailed in 2.1). To prevent data leakage and ensure an unbiased evaluation, we implemented a dataset partition:

- **Development Set** ($N = 40$): Utilized for hyperparameter optimization via 4-fold CV.
- **Held-out Test Set** ($N = 22$): Isolated for final statistical testing.

Baseline Configurations To isolate the specific contributions of our locally-optimized spherical SSL framework, we compare our approach against both spherical and planar baselines. Notably, no large-scale foundation models currently support native spherical topology. We utilize a gnomonic cubemap projection, mapping the sphere onto six cube faces to mitigate the polar singularities inherent to equirectangular projections. In addition to the 3D U-Net volumetric baseline from Borne et al. [1] (the core framework underlying BrainVISA automatic labeling) evaluated on this exact dataset, we define three surface-based baselines to isolate our approach’s specific contributions:

1. Spherical U-Net "From Scratch": We trained the Spherical U-Net architecture. This quantifies the benefit of the pre-learned geometric priors .
2. 2D U-Net (ResNet-18): U-Net architecture with a ResNet-18 backbone.
3. 2D Vision Foundation Models : to adapt DINOv3 [19] for dense sulcal labeling, we employed two decoding strategies. First, we utilize a linear probing approach with unfreezing the final two Transformer blocks. A 1×1 convolution maps geometric inputs to RGB space. As a second strategy, to better capture fine-grained topology, we draw inspiration from the Dense Prediction Transformer (DPT) architecture[16]. Feature representations are extracted from layers 3, 6, 9, and 12 of the backbone network and subsequently fused using a decoder. Implementation details can be found in our code repository.

3.2 Implementation Details and Hyperparameter Optimization

During the SSL phase, models were optimized with a batch size $\in [128, 1024]$, a learning rate $\in [10^{-5}, 10^{-3}]$, and $\lambda \in [0.05, 20]$. The structural augmentations (Sec. 2.2) were bounded as follows: Gaussian spherical rotations ($\sigma \in [2^\circ, 80^\circ]$), patch masking (500 to 15,000 vertices), and vertex dropout (0% to 60%).

Hyperparameters were calibrated via Random Search across a 4-fold cross-validation strictly confined to the development set.

For downstream fine-tuning, the TPI and decoder learning rates were sampled $\in [10^{-6}, 10^{-3}]$ (with the pre-trained encoder’s learning rate strictly being the one-tenth of this value $\in [10^{-7}, 10^{-4}]$ (Sec. 2.3)). Weight decay was bounded $\in [0, 10^{-4}]$. The compound objective balanced a class-weighted Cross-Entropy and a Dice loss (Dice weight $\in [0, 3]$). Spatial augmentations, applied with probability $p \in [0.3, 1.0]$, included feature noise ($\sigma \in [0, 0.4]$) and rotations ($\sigma \in [5^\circ, 90^\circ]$). Empirically, 2D projected baselines performed better without rotational augmentations; so, rotations were set to zero for these specific models. Raw predictions are refined via a mesh-adjacency graph, where isolated connected components smaller than a threshold ($V_{min} = 10$) are reassigned to the mode of their robust topological neighbors.

4 Results

Evaluation metrics: We compute the Error of Sulcal Identification (ESI), defined as a weighted sum of labeling errors across all sulci L : $ESI = \sum_{l \in L} w_l * \frac{FP_l + FN_l}{FP_l + FN_l + 2 * TP_l}$. ESI is intrinsically weighted by sulcal length (wl). Because length definition varies between 3D folds and 2D lines, direct ESI comparisons are irrelevant between them. Furthermore, anatomical analysis is conducted using the localized error rate $E_{local}(l) = \frac{FP_l + FN_l}{FP_l + FN_l + TP_l}$. It reaches 1 when a sulcus is completely hallucinated or missed (frequently for highly variable, minor sulci).

Global Performance and Data Efficiency: As shown in Fig. 2A, our Spherical Transfer framework achieves the highest mean Dice (0.77 ± 0.03) on the held-out test set when trained on the maximum available split ($N = 30$). Our approach visibly reduces performance variance and raises the minimum accuracy




Fig. 2. Quantitative evaluation of the Spherical Transfer Learning framework.

bound, demonstrating enhanced robustness to inter-subject variability. Planar projection strategies, including the vision foundation model DINOv3 (Dense: 0.70 ± 0.05), perform similarly to a standard 2D ResNet-18 (0.69 ± 0.05). DINOv3’s inability to significantly surpass from-scratch baselines is not intended as a claim of absolute superiority, but rather illustrates the structural domain gap between natural RGB images and distorted cortical topologies: even a powerful foundation model fine-tuned via state-of-the-art cubemap projection fails to match a surface-native approach.

The practical value of the geometric SSL priors is most prominent in extreme low-data configurations (Fig. 2B). All training configurations (from 5 to 30 subjects) are strictly evaluated on the same fixed test set ($N=22$).


When the annotated cohort is reduced to merely 5 subjects, the from-scratch spherical baseline exhibits severe topological degradation, with its ESI deteriorating to 0.32 ± 0.06 . In contrast, our transferred model maintains stable identification score (ESI: 0.24 ± 0.045) (Wilcoxon signed-rank test, $p < 0.01$).

Ablation Study: Disentangling Pre-training Strategies: Table 1 isolates the specific contributions of our pre-training and transfer mechanisms. The most critical observation lies in the performance of the frozen locally-optimized encoder. Without updating any backbone weights during the downstream task, it achieves a mean Dice of 0.71 ± 0.04 , performing on par with the fully trained from-scratch baseline (0.72 ± 0.05 , $p > 0.05$). This confirms that the localized SSL pretext task successfully disentangles fine-grained cortical geometry, encoding intrinsic topological boundaries without any semantic supervision. Conversely, a frozen backbone pre-trained on macroscopic targets (Age/Sex) de-

Table 1. Ablation study evaluating the impact of pre-training objectives and transfer initialization on downstream labeling performance.

Pre-training Strategy	Encoder Status	Dice ($\uparrow$)	ESI ($\downarrow$)
None ("From Scratch")	Unfrozen	0.72 ± 0.05	0.21 ± 0.03
Global Supervised (Age/Sex)	Frozen	0.58 ± 0.07	0.32 ± 0.05
Global Supervised (Age/Sex)	Unfrozen	0.70 ± 0.05	0.22 ± 0.04
Locally-Optimized SSL	Frozen	0.71 ± 0.04	0.20 ± 0.04
Locally-Optimized SSL (Random Init)	Unfrozen	0.74 ± 0.03	0.17 ± 0.03
Locally-Optimized SSL	Unfrozen	0.77 ± 0.03	0.15 ± 0.02

grades severely (0.58 ± 0.07), proving that global optimization discards the high-frequency structural gradients essential for dense segmentation. While the frozen local encoder provides a robust geometric foundation, reaching peak performance (Dice: 0.77 ± 0.03 , ESI: 0.15 ± 0.02) requires unfreezing the backbone (with differential learning rates) and adding our soft-initialized TPI. This full transfer framework significantly outperforms the from-scratch baseline ($p < 0.01$).

**Fig. 3.** Surface map of local error variation (ΔE_{local}) against the from-scratch model.

Local Error Analysis: Resolving Topological Ambiguity: Figure 3 maps the local error variation (ΔE_{local}) across the sulci, where negative values denote an error reduction relative to the from-scratch baseline. While our transfer framework improves the delineation of certain large, prominent folds (e.g., Occipital Lobe : -12% , Superior Frontal Sulcus: -6.5%), the largest error reductions are concentrated on some small, highly variable structures. For instance, the Superior Postcentral (-14.82%), Rhinal (-13.94%), and Superior Precentral (-12.23%) sulci exhibit significant performance leaps ($p < 0.05$). By encoding intrinsic curvature and depth variations during pre-training, the network acquires the necessary topological constraints to accurately identify ambiguous folds that are otherwise systematically smoothed over or missed in data-scarce regimes.

5 Conclusion

In this work, we introduced a Spherical Transfer Learning framework to address the challenge of exhaustive cortical sulci identification in data-scarce regimes.

Our results demonstrate that decoupling geometric representation learning from semantic fine-tuning (achieved via a locally-optimized SSL prior and a soft-initialized TPI) effectively mitigates the statistical regularity bias inherent to standard supervised models. Notably, this geometric disentanglement yields the most significant identification error reductions on highly variable, tertiary sulci.

Rather than relying on architectural complexity, this study demonstrates the fundamental necessity of localized geometric priors for surface-native neuroanatomy. Future work will focus on scaling these localized pretext tasks to higher-capacity paradigms to further advance dense cortical mapping.

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), by AXA foundation for the project premaIA, by the "Fondation de France" for the project TSA-ideogrammes, by the PEPR Digital Health via BHT (ANR-22-PESN-0012) and the Fondation Paralysie Cerebrale to the ENSEMBLE project. This project was provided with computing HPC and storage resources by **GENCI TGCC**, thanks to the grant **2024-A0170313800** on the **Joliot Curie's ROME** supercomputer partition. This project was provided with computer and storage resources by **GENCI at IDRIS** thanks to the grant 010316018 on the supercomputer Jean Zay's H100 partition.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

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 (2020)
2. 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)
3. Cohen, T.S., Geiger, M., Köhler, J., Welling, M.: Spherical CNNs (Feb 2018)
4. Cox, J., Liu, P., Stolte, S.E., Yang, Y., Liu, K., See, K.B., Ju, H., Fang, R.: Brain-SegFounder: Towards 3D Foundation Models for Neuroimage Segmentation (Nov 2024), arXiv:2406.10395 [eess]
5. Dale, A.M., Fischl, B., Sereno, M.I.: Cortical Surface-Based Analysis: I. Segmentation and Surface Reconstruction. *NeuroImage* **9**(2), 179–194 (Feb 1999)
6. Dufournet, A., Laval, J., Rivière, D., de Vareilles, H., Murray, G.K., Cachia, A., Chavas, J., Frouin, V., Mangin, J.F.: A self-supervised learning framework for discovering cortical folding patterns under genetic influence: Application to the Anterior Cingulate Cortex. *Imaging Neuroscience* **3**, IMAG.a.987 (Nov 2025)
7. Fischl, B., Sereno, M.I., Dale, A.M.: Cortical Surface-Based Analysis: II: Inflation, Flattening, and a Surface-Based Coordinate System. *NeuroImage* **9**(2), 195–207 (Feb 1999)
8. Hao, L., Bao, S., Tang, Y., Gao, R., Parvathaneni, P., Miller, J.A., Voorhies, W., Yao, J., Bunge, S.A., Weiner, K.S., Landman, B.A., Lyu, I.: Automatic Labeling

- of Cortical Sulci Using Spherical Convolutional Neural Networks in a Developmental Cohort. In: 2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI). pp. 412–415. IEEE, Iowa City, IA, USA (Apr 2020)
9. Kondepudi, A., Rao, A., Zhao, C., Lyu, Y., Harake, S., Banerjee, S., Joshi, R., Meissner, A.K., Hou, R., Jiang, C., Chowdury, A., Srinivasan, A., Athey, B., Gulani, V., Pandey, A., Lee, H., Hollon, T.: Health system learning achieves generalist neuroimaging models (Nov 2025), arXiv:2511.18640 [cs]
 10. Laval, J., Chavas, J., Troiani, V., Snyder, W., Patti, M., Moyal, M., Plaze, M., Cachia, A., Yi Sun, Z., Frouin, V., Gori, P., Rivière, D., Mangin, J.F.: Towards a Foundation Model for Cortical Folding. In: Machine Learning in Clinical Neuroimaging: 7th International Workshop, MLCN 2024, Held in Conjunction with MICCAI 2024, Marrakesh, Morocco, October 10, 2024, Proceedings. pp. 78–88. Springer-Verlag, Berlin, Heidelberg (Dec 2024)
 11. Lyu, I., Bao, S., Hao, L., Yao, J., Miller, J.A., Voorhies, W., Taylor, W.D., Bunge, S.A., Weiner, K.S., Landman, B.A.: Labeling lateral prefrontal sulci using spherical data augmentation and context-aware training. *NeuroImage* **229**, 117758 (Apr 2021)
 12. Lyu, I., Kim, S.H., Woodward, N.D., Styner, M.A., Landman, B.A.: TRACE: A Topological Graph Representation for Automatic Sulcal Curve Extraction. *IEEE Transactions on Medical Imaging* **37**(7), 1653–1663 (Jul 2018)
 13. Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., Bronstein, M.M.: Geometric Deep Learning on Graphs and Manifolds Using Mixture Model CNNs. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5425–5434. IEEE, Honolulu, HI (Jul 2017)
 14. Parvathaneni, P., Nath, V., McHugo, M., Huo, Y., Resnick, S.M., Woodward, N.D., Landman, B.A., Lyu, I.: Improving human cortical sulcal curve labeling in large scale cross-sectional MRI using deep neural networks. *Journal of Neuroscience Methods* **324**, 108311 (Aug 2019)
 15. Perrot, M., Rivière, D., Mangin, J.F.: Cortical sulci recognition and spatial normalization. *Medical Image Analysis* **15**(4), 529–550 (Aug 2011)
 16. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision Transformers for Dense Prediction (Mar 2021), arXiv:2103.13413 [cs]
 17. Rivière, D., Geffroy, D., Denghien, I., Souedet, N., Cointepas, Y.: BrainVISA: an extensible software environment for sharing multimodal neuroimaging data and processing tools. *NeuroImage* **47**, S163 (Jul 2009)
 18. Robinson, E.C., Garcia, K., Glasser, M.F., Chen, Z., Coalson, T.S., Makropoulos, A., Bozek, J., Wright, R., Schuh, A., Webster, M., Hutter, J., Price, A., Cordero Grande, L., Hughes, E., Tusor, N., Bayly, P.V., Van Essen, D.C., Smith, S.M., Edwards, A.D., Hajnal, J., Jenkinson, M., Glocker, B., Rueckert, D.: Multimodal surface matching with higher-order smoothness constraints. *NeuroImage* **167**, 453–465 (Feb 2018)
 19. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., Bojanowski, P.: DINOv3 (Aug 2025), arXiv:2508.10104 [cs]
 20. Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M., Liu, B., Matthews, P., Ong, G., Pell, J., Silman, A., Young, A., Sprosen, T., Peakman, T., Collins, R.: UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS medicine* **12**(3), e1001779 (Mar 2015)

21. 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) (Jan 2018), publisher: Society for Neuroscience Section: New Research
22. Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow Twins: Self-Supervised Learning via Redundancy Reduction (Jun 2021), arXiv:2103.03230 [cs]
23. Zhao, F., Xia, S., Wu, Z., Duan, D., Wang, L., Lin, W., Gilmore, J.H., Shen, D., Li, G.: Spherical U-Net on Cortical Surfaces: Methods and Applications. *Information processing in medical imaging : proceedings of the ... conference* **11492**, 855–866 (Jun 2019)