Segmentation of supragranular and infragranular layers in ultra-high resolution 7T *ex vivo* MRI of the human cerebral cortex

Xiangrui Zeng[1,2], Oula Puonti[1,3], Areej Sayeed[1,2], Rogeny Herisse[1,2], Jocelyn Mora[1,2], Kathryn Evancic[1,2], Divya Varadarajan[1,2], Yael Balbastre[1,2], Irene Costantini[4,5,6], Marina Scardigli[5], Josephine Ramazzotti[5], Danila DiMeo[5], Giacomo Mazzamuto[4,5,7], Luca Pesce[5], Niamh Brady[5], Franco Cheli[5], Francesco Saverio Pavone[4,5,7], Patrick R. Hof[8], Robert Frost[1,2], Jean Augustinack[1,2], André van der Kouwe[1,2], Juan Eugenio Iglesias[1,2,*], and Bruce Fischl[1,2,*]

[1]Department of Radiology, Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA, USA

[2]Harvard Medical School, Department of Radiology, Boston, MA, USA

[3]Danish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital - Amager and Hvidovre, Copenhagen, Denmark

[4]National Research Council - National Institute of Optics (CNR-INO), Sesto Fiorentino, Italy

[5]European Laboratory for Non-Linear Spectroscopy (LENS), Sesto Fiorentino, Italy

[6]Department of Biology, University of Florence, Italy

[7]Department of Physics and Astronomy, University of Florence, Italy

[8]Nash Family Department of Neuroscience and Friedman Brain Institute, Icahn School of Medicine at Mount Sinai, New York, NY, USA

[9]Centre for Medical Image Computing, University College London, London, UK

[10]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA

*Corresponding authors. Email address: jiglesiasgonzalez@mgh.harvard.edu, fischl@nmr.mgh.harvard.edu

Abstract

Accurate labeling of specific layers in the human cerebral cortex is crucial for advancing our understanding of neurodevelopmental and neurodegenerative disorders. Leveraging recent advancements in ultra-high resolution *ex vivo* MRI, we present a novel semi-supervised segmentation model capable of identifying supragranular and infragranular layers in *ex vivo* MRI with unprecedented precision. On a dataset consisting of 17 whole-hemisphere *ex vivo* scans at 120 μm, we propose a multi-resolution U-Nets framework (MUS) that integrates global and local structural information, achieving reliable segmentation maps of the entire hemisphere, with Dice scores over 0.8 for supra- and infragranular layers. This enables surface modeling, atlas construction, anomaly detection in disease states, and cross-modality validation, while also paving the way for finer layer segmentation. Our approach offers a powerful tool for comprehensive neuroanatomical investigations and holds promise for advancing our mechanistic understanding of progression of neurodegenerative diseases.

Keywords: *Ex vivo* MRI; Cortical Layers; Segmentation; and High Resolution

1 Introduction

The human neocortex is a complex structure organized into a number of distinct layers, characterized by variations in the size and packing density of their constituent neurons. These layers form during cortical development as a result of radial and tangential neuronal migration [1]. During embryonic development, newly generated neocortical projection neurons migrate along radial glia in successive waves, leading to the formation of cortical layers in an inside-out pattern [2]. This means that the deepest layers are populated first, while the most superficial layers are occupied by the last-generated neurons. In addition to their unique organization, cortical layers also exhibit distinct patterns of connectivity. For example, pyramidal neurons in layers II and III predominantly project to other cortical regions, while those in layer V project mainly to the striatum and brainstem, and those in layer VI project to the thalamus [3].

When it comes to diseases affecting the human neocortex, specific layers or cell types often show particular pathologies. For instance, in schizophrenia, large pyramidal cells in layer III display reduced cell size [4]. Deficits in reelin expression are primarily found in the most superficial layers, also known as supragranular layers, among schizophrenia [5], bipolar disorder [6] and autism spectrum disorder [7] patients. Another example is the development of Alzheimer's disease pathology, which includes the neuronal loss within the neocortex, and initially manifests in the superficial cortical layers (II–IV) during its early stages. As

the disease progresses, it extends to affect the deeper layers (V–VI) [8]. These examples highlight the importance of correctly annotating specific layers in the human neocortex. This identification is essential for advancing our understanding of these disorders and may provide valuable insights for potential therapeutic approaches.

A natural approach to identifying cortical layers is microscopic examination of tissue histology. While histology offers definitive insights into microscopic tissue morphology, it suffers from limitations such as sampling bias due to a restricted field of view – and therefore has difficulties in exploring the interrelationships between different and potentially dysfunctional regions [9]. Moreover, histology is labor-intensive and invasive, which may decrease the measured cortical thickness by factors like dehydration and increase by factors like the slicing direction [10, 11]. In general, any 2D technique for measuring cortical properties suffers inaccuracies that arise due to the effects of through-plane folding. In the context of cortical thickness, any 2D measure will inevitably overestimate it except in locations where the cut is perfectly orthogonal to the cortex.

In contrast, conventional *in vivo* MRI can provide isotropic whole-brain images rapidly and non-invasively, but lacks the resolution and specificity of histology [12]. Unlike *in vivo* MRI, *ex vivo* MRI is not affected by motion artifacts and has much less limiting time constraints [13]. Extended scanning time enables increased spatial resolution, which in turn enables visualization of mesoscale neuroanatomy. The resulting increase in spatial resolution is crucial for visualizing mesoscale neuroanatomy, such as cortical layers and subcortical nuclei, which are challenging to visualize in even the highest-resolution *in vivo* MRI datasets [14]. *Ex vivo* MRI also circumvents the spatial distortions (tearing or folding) associated with histological methods during brain tissue fixation, embedding, sectioning, and slide-mounting [15]. This makes it well-suited for characterizing neuroanatomy at high resolution and providing finer macroscopic morphometric measures, such as cortical thickness, of the underlying cytoarchitecture. Although imaging the intact human brain *ex vivo* at high magnetic fields is challenging due to the need for specialized hardware [13], recent progress in high-field scanner and coil technology, and imaging protocols [16], has enabled full-brain scanning with voxel sizes as small as 100 μm [13, 17], helping bridge the gap between histology and MRI.

Equipped with these imaging advances [18], we now have the means to acquire data sets consisting of high-resolution, whole-hemisphere scans from multiple post-mortem subjects. Previous high-resolution data sets, such as BigBrain [19] and the Allen Brain Atlas [20], include only a single human brain with whole-brain MRI and histology, which prevents us from reliably quantifying the inter-subject variability of human neuroanatomy. However, manually labeling large multi-subject data sets is not feasible in practice and existing automated tools for segmenting the supra- and infragranular layers require a large amount of

3

manually prepared training data [21]. In general, automated segmentation of *ex vivo* data is hindered by limited training data, and the few existing data sets that include multiple subjects only cover specific sub-structures [22, 23].

Convolutional Neural Networks (CNNs) are becoming increasingly popular in medical image analysis [24]. When training data were available, processing large 3D volumes using CNNs is challenging due to limitations in Graphics Processing Unit (GPU) memory. Downsampling the volumes to reduce the memory load will inevitably lead to a loss of fine structural details, resulting in decreased segmentation accuracy. Similarly, using subvolume patches can result in reduced accuracy due to the lack of global context information. To address these issues, researchers have proposed 2.5D segmentation approaches that operate on orthogonal planes and subsequently merge their information [25]; CNNs with separate high- and low-resolution paths [26]; and light-weight models using dilated convolution [27] or group normalization [28]. Isensee et al. [29] have developed the nnU-Net ("no-new-Unet") framework as a robust and self-adjusting extension of the U-Net. This framework involves minor modifications to both the 2D and 3D U-Net designs, wherein 2D and 3D connections are integrated to collaboratively establish a network pool. Nevertheless, there are no existing methods with end-to-end training that can effectively incorporate important 3D global context information with local high-resolution details needed for accurate labeling of the supra- and infragranular layers.

In this paper, we present a dataset [12] consisting of 17 whole-hemisphere *ex vivo* scans at 120 μm with partial manual annotations and propose a semi-supervised model, Multi-resolution UNets Semi-supervised (MUS), that require a minimal amount of annotated training data to segment supra- and infragranular layers in ultra-high resolution *ex vivo* brain MRI. A variant of the U-Net [30], the multi-resolution U-Nets is designed to incorporate both global and local structural information for high-resolution segmentation accuracy. With this segmentation model, we obtained, for the first time, reliable segmentation maps (Dice score: >0.8) of supra- and infragranular layers over the whole hemisphere. The combination of the unique dataset and novel automated segmentation approach, paves the way for an in-depth examination of cortical layer organizations and will allow us to (1) place surface models and build atlases; (2) infer laminar anomalies between disease stages and healthy controls using the atlases; (3) benchmark and validate cortical layer segmentation results in other imaging modalities; and (4) progress to finer segmentation of more cortical layers in the future. The dataset and segmentations can be downloaded from the DANDI data archive[1], and the code will be distributed with the FreeSurfer[2] software suite.

---

[1]https://www.dandiarchive.org
[2]https://surfer.nmr.mgh.harvard.edu

2 Materials and methods

2.1 Datasets

MRI scans of 17 whole hemispheres (see demographics in Table 1) were acquired on a Siemens 7 Tesla scanner using a custom built 32-channel receive array as detailed in [13]. The scans were acquired at 120 μm isotropic resolution (Figure 1), which allows reliable visual identification of the supra- and infragranular boundary throughout the neocortex. This was achieved by using a multi-echo spoiled gradient echo sequence (ME-GRE) and acquired a series of images at different flip angles [31]. The k-space acquisition was segmented to fit data from a single segment into scanner memory, then streamed to a dedicated computer for offline reconstruction. Adjacent k-space segments were modified to contain a small number of overlapping lines enabling us to correct for phase discontinuities due to field drift during the extremely long scans (14 hours per volume). The scans are also sensitive to various distortions and intensity inhomogeneities due to variations in B0 and B1-/+ fields, which we mapped and corrected following the procedures in [12]. Briefly, the alternating reversing-polarity reads of the multi-echo fast low angle shot (MEF) scans provide a mechanism for mapping and correcting B0 field inhomogeneities at the intrinsic resolution of the scans when combined with a low-resolution field map [32, 33]. For transmit (B1+) inhomogeneities, we acquired several low-resolution scans with varying transmit voltage to map the flip angle field, then used these maps and the B0-corrected MEF scans as inputs to the steady-state MEF equations [31], yielding a set of synthesized scans that are of higher SNR than the individual input scans. Finally, we used the SAMSEG algorithm from Freesurfer [34] to correct for receive inhomogeneities (B1-).

Table 1: Demographic information of donor cohorts in the *ex vivo* MRI dataset. Interval denotes post-mortem interval in hours.

Donor cohort	17 (Right: 6, Left: 10, whole-brain: 1)	
Sex	Male: 11	Female: 6
Age	64.5 ± 11.2	66.3 ± 6.7
Interval	19.4 ± 5.9	15.8 ± 8.6

An relevant question is what does the visible contrast boundary in the cortical gray matter represent (Figure 2(B)). Previously, we have performed confocal light-sheet fluorescence microscopy (LSFM) [12] on tissue slabs from BA 44/45 treated with the SHORT clearing technique [35]. Supra- and infragranular labels can be derived from LSFM as neuron subtypes are specifically labeled. Using LSFM registered into the MRI space, we demonstrate

that the MR contrast boundary in cortical gray matter corresponds to the cytoarchitectural boundary between layers III and V, visible on the NeuN stain (Figure 2(A)). Based on the myelin density differences [36, 37] and their resulting contrast in MRI, we group layer I, II, and III together as the supragranular layer, and layer IV (absent in some regions), V, and VI together as the infragranular layer.
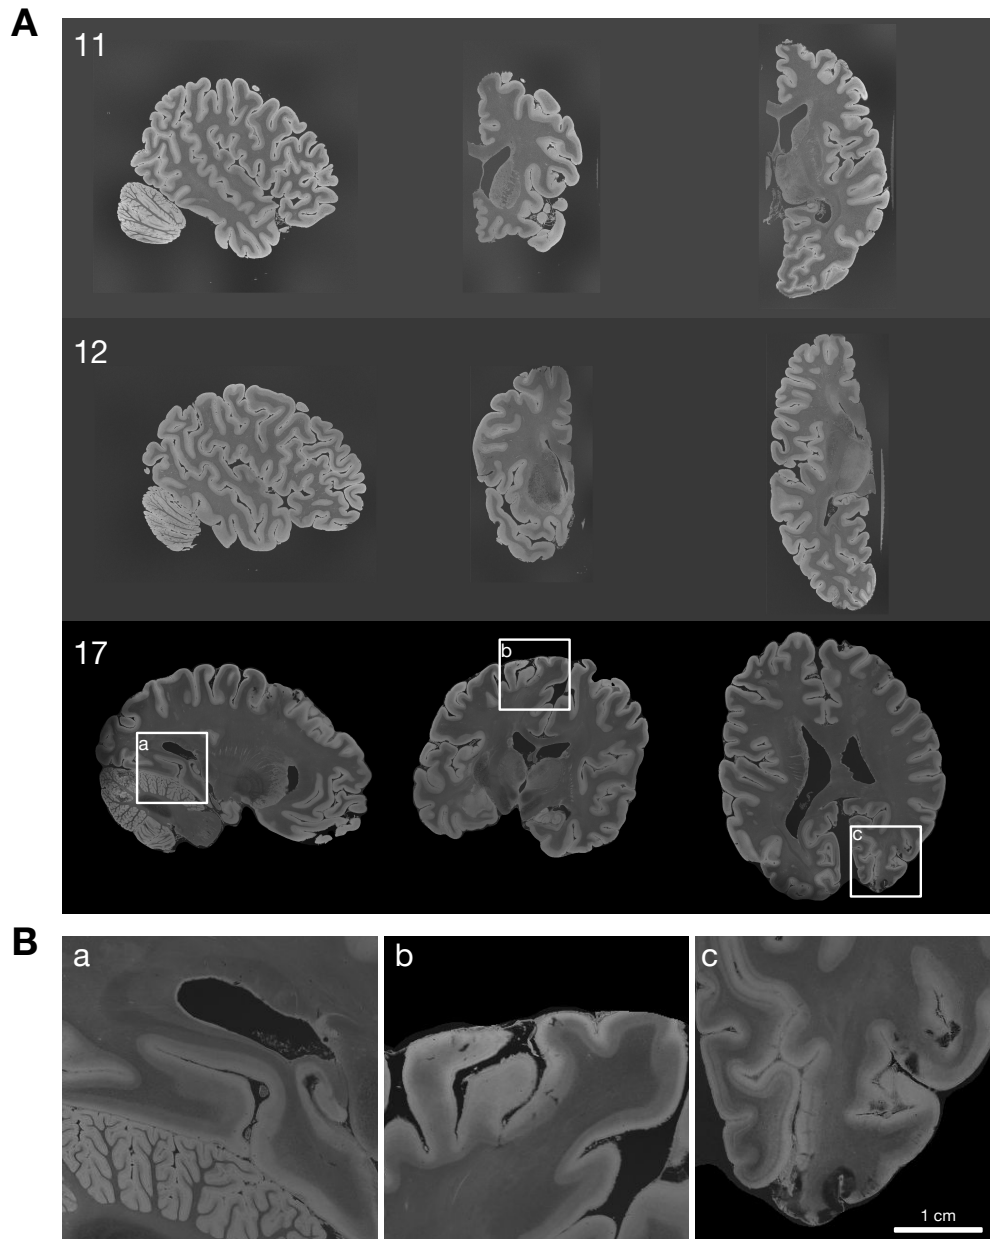

Figure 1: **A.** Sagittal (left), coronal (center), and axial (left) slices of case 11, 12, and 17. Bias correction was applied for all scans. Comprehensive visualization of all cases can be found in the supplementary material. **B.** Zoomed views of selected regions in the whole brain case 17.

2.2 Data preprocessing

Manual annotation was performed through *Freeview* tool in *FreeSurfer* [38] to label the visible supra-/infra-laminar boundary. The supragranular layers appear as a bright band in the neocortex, while the infragranular layers appear as a slightly darker band. The white matter appears as the dark area interior to the neocortex. Using these intensity characteristics, we manually segmented these three structures and the background from 100 slices in the Brodmann area (BA) 44/45 of each hemisphere specimen using the coronal view on *Freeview*. We maintained this single plane for manual labeling to limit bias in the labeling process, while inspecting other planes to avoid jagged reconstruction. In addition, we labeled two samples across the whole-hemisphere, one slice every 40. In total, about 3% voxels of supra- and infragranular layers are manually labeled. Figure 2(B) shows examples of manual annotation on selected samples.

An additional background (not cerebral gray or white matter) labeling for training the MUS segmentation model was created using a combination of the segmentation outputs from SynthSeg [39] and SAMSEG [34]. Specifically, the *ex vivo* scans were first downsampled to 500 μm isotropic resolution, then processed with both SynthSeg and SAMSEG to produce probabilistic structure maps with values ranging from 0 to 1, and finally these maps were combined to produce a single background probability map consisting of all structures except the cortical gray matter. These maps were then upsampled back to the original resolution and thresholded at a value of 1.0 to get the final background labeling.

A
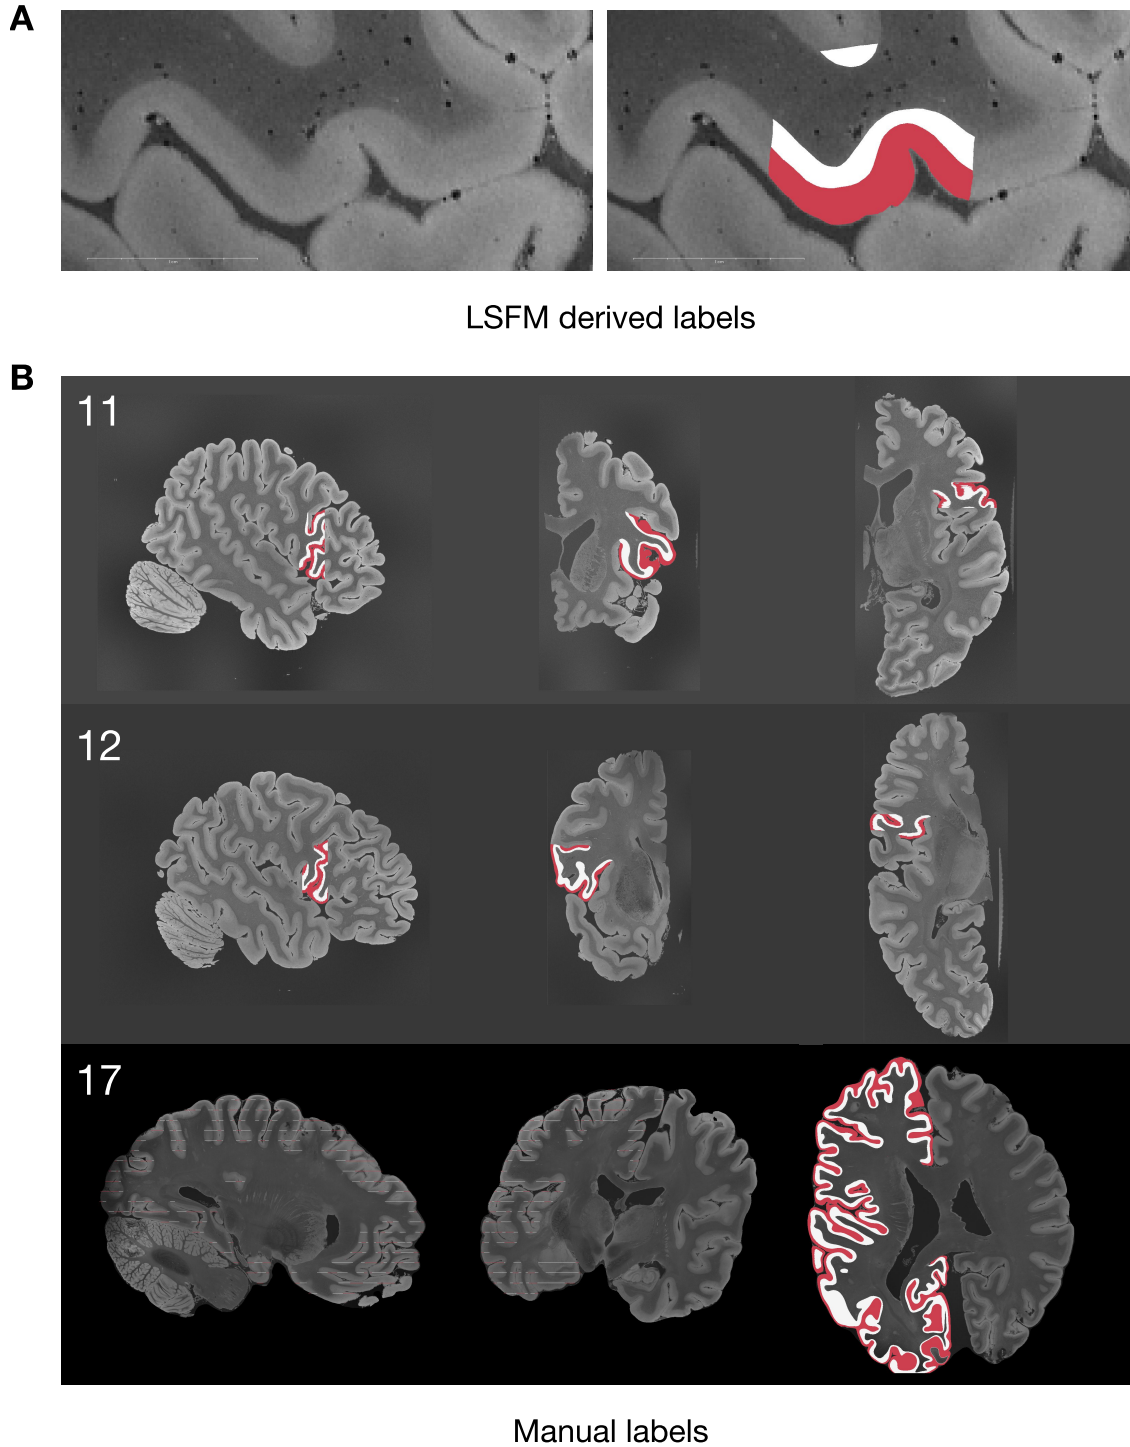

LSFM derived labels

B

Manual labels

Figure 2: **A.** LSFM derived supra- and infragranular layer labels co-registered with MRI [12]. **B.** 2D slices of case 11, 12, 17 at sagittal, coronal, and axial view with manual annotation overlaid (red: supragranular layer, white: infragranular layer). Comprehensive visualization of annotations on all samples can be found in the supplement.

2.3   Semi-supervised segmentation model

The U-Net is a deep learning architecture that has gained significant attention and popularity within the field of medical imaging. Initially proposed in 2015 by [30], it resembles an autoencoder, in the sense that it consists of a contracting path that captures semantic context and a symmetric expansive path that enables precise feature localization. Crucially, encoder features are concatenated with features at the same resolution level in the decoder via skip connections, which effectively preserve high-frequency components of the signal that enable segmentation of convoluted boundaries. This design facilitates the incorporation of both global and local information, making it particularly effective for tasks where accurate delineation of structures is crucial, such as in identifying organs [40, 41], tumors [42], and anatomical features [43, 39, 44]. Beyond image segmentation, the U-Net's versatility has led to its adoption in various medical imaging applications, including image denoising [45, 46], registration [47, 48], and super-resolution [49, 50, 51], showcasing its adaptability and robust performance across different scenarios.

Multi-resolution U-Nets: To overcome the limitations related to the size of the data set and sparse annotations described in the introduction, we propose a cascaded resolution approach, inspired by previous works [29, 52], in combination with semi-supervised learning, which takes in volumetric inputs downsampled at different resolutions, while ensuring that all U-Net components receive inputs of the same size. This enables us to simultaneously capture both a large field of view and fine structural details. Our multi-resolution U-Net architecture is depicted in Figure 3(A) and employs a series of cascaded U-Net components. At a coarse resolution, the U-Net input volumes have a larger field of view but lack fine structural details. Conversely, at a fine resolution, the field of view is smaller, but fine structural details are preserved. By utilizing features extracted from highly downsampled volumes, we capture global context information, which is then integrated with features from volumes of the original resolution. Each component U-Net follows a standard U-Net architecture (Figure 3(B)). During the forward pass, features from the corresponding volume are extracted from the penultimate layer of the U-Net and concatenated to the second layer of the next U-Net. This process ensures the incorporation of spatially matched information from different scales to improve the overall segmentation accuracy.
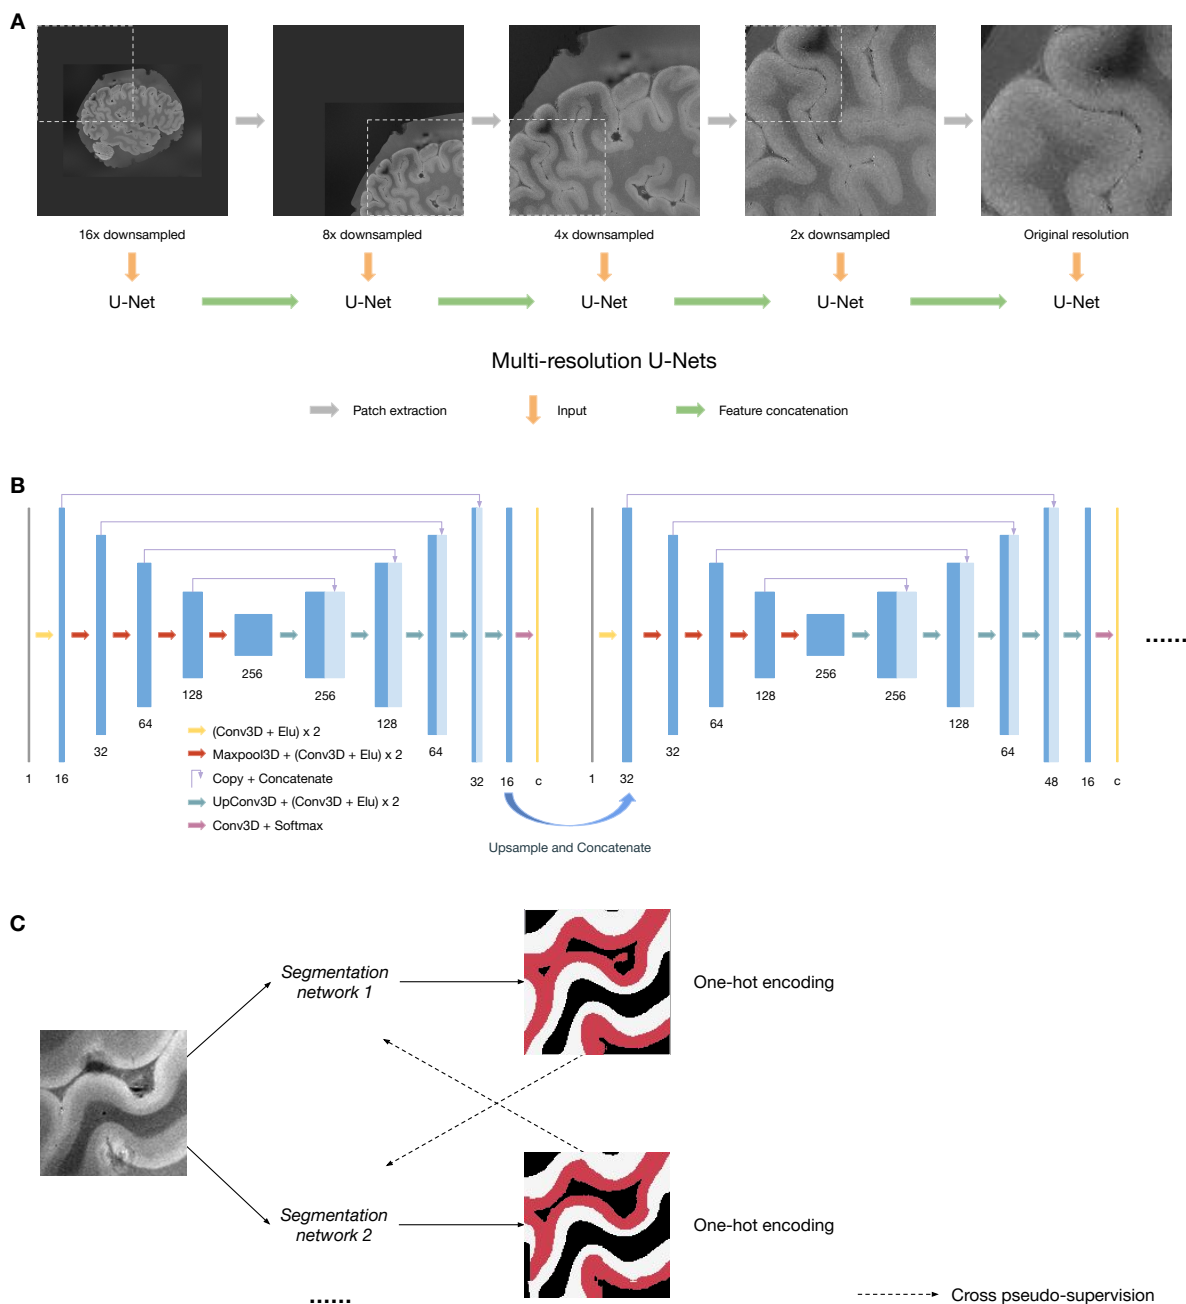
9

Figure 3: **A.** Processing large *ex vivo* MRI volumes using multi-resolution U-Nets. Inputs are downsampled at different scales mimicking a zoom-in procedure. Features extracted from coarser resolutions contain global context information and are integrated into subsequent U-Nets. **B.** Model architecture of component U-Nets. Features from the second layers are extracted, upsampled, and concatenated to the second layer of the next U-Net. All component U-Nets are trained simultaneously in an end-to-end fashion. **C.** Cross pseudo supervision is a semi-supervised learning technique that trains two or more networks at the same time and uses their outputs to supervise each other.

10

For the task of automatically segmenting supra- and infragranular layers in the *ex vivo* MRI dataset, we would ideally have a number of hemisphere samples fully labeled manually. This manual segmentation could then be used to train CNNs in a supervised fashion, in order to automatically predict labels on other samples by mimicking the manual segmentation procedure. However, 3D ultra-high resolution *ex vivo* MRI data is very large and thus extremely time-consuming and laborious to manually annotate. In semi-supervised training, the network learns from both labeled and unlabeled data to train a predictive model; the latter is often relatively easier to obtain in much larger amounts. Semi-supervised training of CNNs mainly relies on the idea of incorporating knowledge priors [53, 54] or enforcing consistency between labeled ground truth and predictions from unlabeled data [55, 56]. Here, we propose a semi-supervised training strategy to effectively utilize the large amount of unlabeled data to improve the segmentation performance. Our semi-supervised segmentation approach is mainly adapted from the so-called cross pseudo-supervision strategy [57].

Semi-supervised segmentation: from a set of MRI volumes $x \in X$, we aim to predict one-hot segmentation $y \in Y$. We denote labeled MRI volume as $x^l \in X^l$ with segmentation labels $y^l \in Y^l$ and unlabeled MRI volumes as $x^u \in X^u$. Two segmentation networks with identical architecture f_{θ_1} and f_{θ_2} are initialized with different random weights. These two CNNs are trained with two loss functions defined symmetrically. In regions with existing manual segmentation labels, we directly compare the network prediction using one-hot encodings of the ground truth:

$$L_l = \text{Dice}(f_{\theta_1}(x^l), y^l), \tag{1}$$

where $\text{Dice}(,)$ denotes the soft Dice loss function [58]:

$$\text{Dice}(Y, \hat{Y}) = 1 - \frac{1}{K} \sum_{k=1}^{K} \frac{2 \times \sum_{x,y,z} Y_k(x,y,z)\hat{Y}_k(x,y,z)}{\sum_{x,y,z} Y_k(x,y,z)^2 + \hat{Y}_k(x,y,z)^2}. \tag{2}$$

The vast majority of regions in the training data are unlabeled. In order to utilize the large amount of unlabeled data for improving the segmentation performance, we adapt a cross pseudo-supervision loss function on unlabeled data [57]:

$$L_u = \text{Dice}(f_{\theta_1}(x^u), 1_c \circ f_{\theta_2}(x^u)), \tag{3}$$

where 1_c denotes the one-hot encoding function. The benefits of this approach are twofold. First, it promotes consistent predictions across differently initialized networks for the same input image, improving reliability and decision boundary placement. Second, during later optimization stages, the pseudo labeled data acts as an expansion of the training dataset to

enhance training as compared to using labeled data alone.

Since the segmentation network operates at multiple resolutions, we also enforce the prediction to be consistent at different resolutions using a multi-resolution consistency loss:

$$L_c = \text{Dice}(f_{\theta_1}^{\downarrow}(x^{\downarrow}), 1_c \circ f_{\theta_2}(x)^{\downarrow}), \qquad (4)$$

where $\downarrow$ denotes the downsampling operator.

One potential issue of cross pseudo supervision is error accumulation: In the late training stage, the predictions from the two networks will converge and may be trapped in local optima because errors will also be mutually learned and reinforced. One way to address this issue is to encourage the errors made by the two networks during training to be diverse. We therefore design an error diversity loss function based on the idea that on the labeled region, if both networks make incorrect prediction as compared to the ground truth, we encourage them to make different errors:

$$L_e = 1_{\{y^l \neq 1_c \circ f_{\theta_1}(x^l) \cap y^l \neq 1_c \circ f_{\theta_2}(x^l)\}} \cdot -\text{Dice}(f_{\theta_1}(x^l), f_{\theta_2}(x^l)), \qquad (5)$$

where $1\{\}$ denotes the indicator function.

2.4   Implementation details

All CNN models are implemented using the *PyTorch* framework [59]. All scans were corrected for biases. Each input supplied to the multi-resolution U-Nets contains 5 volumes at different resolutions, and thus has dimensions of 5×128^3. The image patch cascades undergo successive downsampling, reducing dimensions by 16, 8, 4, 2, and 1-fold along all three axes. Consequently, the segmentation output maintains the same hierarchical structure with dimensions of 5×128^3, where each voxel corresponds to a semantic class label.

In the context of labeled input data, a crucial distinction is made based on whether the fifth input volume (original resolution, no downsampling) contains manually labeled supra- or infragranular layer class voxels. Inputs meeting this criterion are categorized as labeled, while others are categorized as unlabeled. During the initial training phase, when model predictions lack precision to mutually guide one another, a strategic approach is employed to progressively enhance the influence of unlabeled samples as training advances. To this end, a parameter ϵ is defined as $\epsilon = e^{\frac{-5(M-m)}{M}}$, where m denotes the current epoch number and M represents the total number of epochs. In each epoch, for every labeled input, the loss is computed as $L = \lambda_1 L_l + \lambda_2 L_e$. For unlabeled inputs, there exists an ϵ probability of being chosen for training, with the loss calculated as $L = \epsilon(\lambda_3 L_u + \lambda_4 L_c)$. The choice of

loss weights is empirical: $\lambda_1 = 1$, $\lambda_2 = 0.1$, $\lambda_3 = 1$, and $\lambda_4 = 0.1$. Since the range of all loss functions is between 0 and 1, these values strike a balance between the importance of different loss functions.

Training spans 1000 epochs, with each epoch involving the loading of inputs from a randomly chosen scan into the training dataset. The input was augmented with random bias field, γ transformation, and Gaussian noise. A batch size of 1 is set, and the training employs the *Adam* optimizer [60] with a learning rate of 1^{-4}. In contrast to training only two networks, our modified cross pseudo-supervision approach involves training three networks to maintain a "backup", thereby enhancing training stability and overall performance. At each step, the two networks with the most dissimilar segmentation predictions, as assessed by the Dice score, are chosen for cross pseudo supervision.

During the prediction stage, an overlapping tile strategy [30] is adopted to ensure smoothness at boundaries.

3 Results

3.1 Evaluation

To assess the performance of automatic supra-/infragranular layer segmentation, we conducted manual segmentation on specific slices. The selection of validation slices followed a structured procedure: 1. Each hemisphere sample underwent parcellation into 14 cortical regions using the *recon-all* tool within *FreeSurfer* [38]. 2. Within each region, a random point was chosen on the white matter surface. 3. The orientation (axial, coronal, or sagittal) most perpendicular to this surface point was determined, and a slice of dimensions 128^2 was extracted centered on this point. 4. Manual segmentation was carried out on the central region of this slice.

In total, 210 slices were chosen for evaluation. To gauge the reliability of the manual segmentation procedure, we randomly picked one slice from each cortical region, which was re-annotated by the same labeler after a 4-week interval. This allowed us to estimate intra-rater variability.

15-fold cross-validation was used in all experiments. In each fold, the 2 samples with sparse whole-hemisphere slice labeling and the 14 samples except the particular one for prediction were used as the training set.

We applied the public implementation of nnU-Net [29] available on GitHub for comparison with our method. Since our training data annotations contain unlabeled parts, we masked out unlabeled parts during the calculation of loss gradient.

3.2 Segmentation map of supragranular and infragranular layers

As a baseline method, we first applied nnU-Net, a widely recognized implementation of
U-nets that provides state-of-the-art results in an array of medical image segmentation
tasks [29]. Given our ultra-high-resolution dataset, nnU-Net self-configured a pipeline with
two training stages. In the first stage, a U-Net was trained on a downsampled version of the
dataset, enabling the entire 3D volume to be processed by the U-Net. In the second stage,
another U-Net was trained on 3D sub-volumes extracted from the whole volume, maintain-
ing full resolution. The sub-volumes, along with their corresponding coarse segmentation
from the first stage, were concatenated as the input. The predictions from this second stage
were kept as the final results. However, as illustrated in Figure 4(A), this approach yielded
suboptimal results, notably missing portions of the neocortex in the layer segmentation.
This failure is likely due to nnU-Net not being tailored for scenarios with limited labeled
training data. Consequently, conventional supervised U-Net models proved insufficient in
achieving our objective of accurately segmenting supra- and infragranular layers under these
constraints.

In contrast to nnU-Net, which exclusively employs labeled data for training, our model
adopts a semi-supervised approach, utilizing both labeled and the substantial majority of
unlabeled data ($\sim 97\%$) for training. Furthermore, while nnU-Net also employs a multi-
resolution strategy, it is limited to two stages and trained independently. In contrast, our
multi-resolution U-Nets can operate across a larger number of resolutions while being trained
in an end-to-end manner, which effectively leverages information from all resolutions and
scales. This approach led to more accurate segmentation of the supra- and infragranular
layers, as shown in Figure 4(A), while excluding non-targeted regions such as the cerebellum.
Additionally, as shown in Figure 4(B), our method qualitatively yields the highest consistency
with the manual annotations.
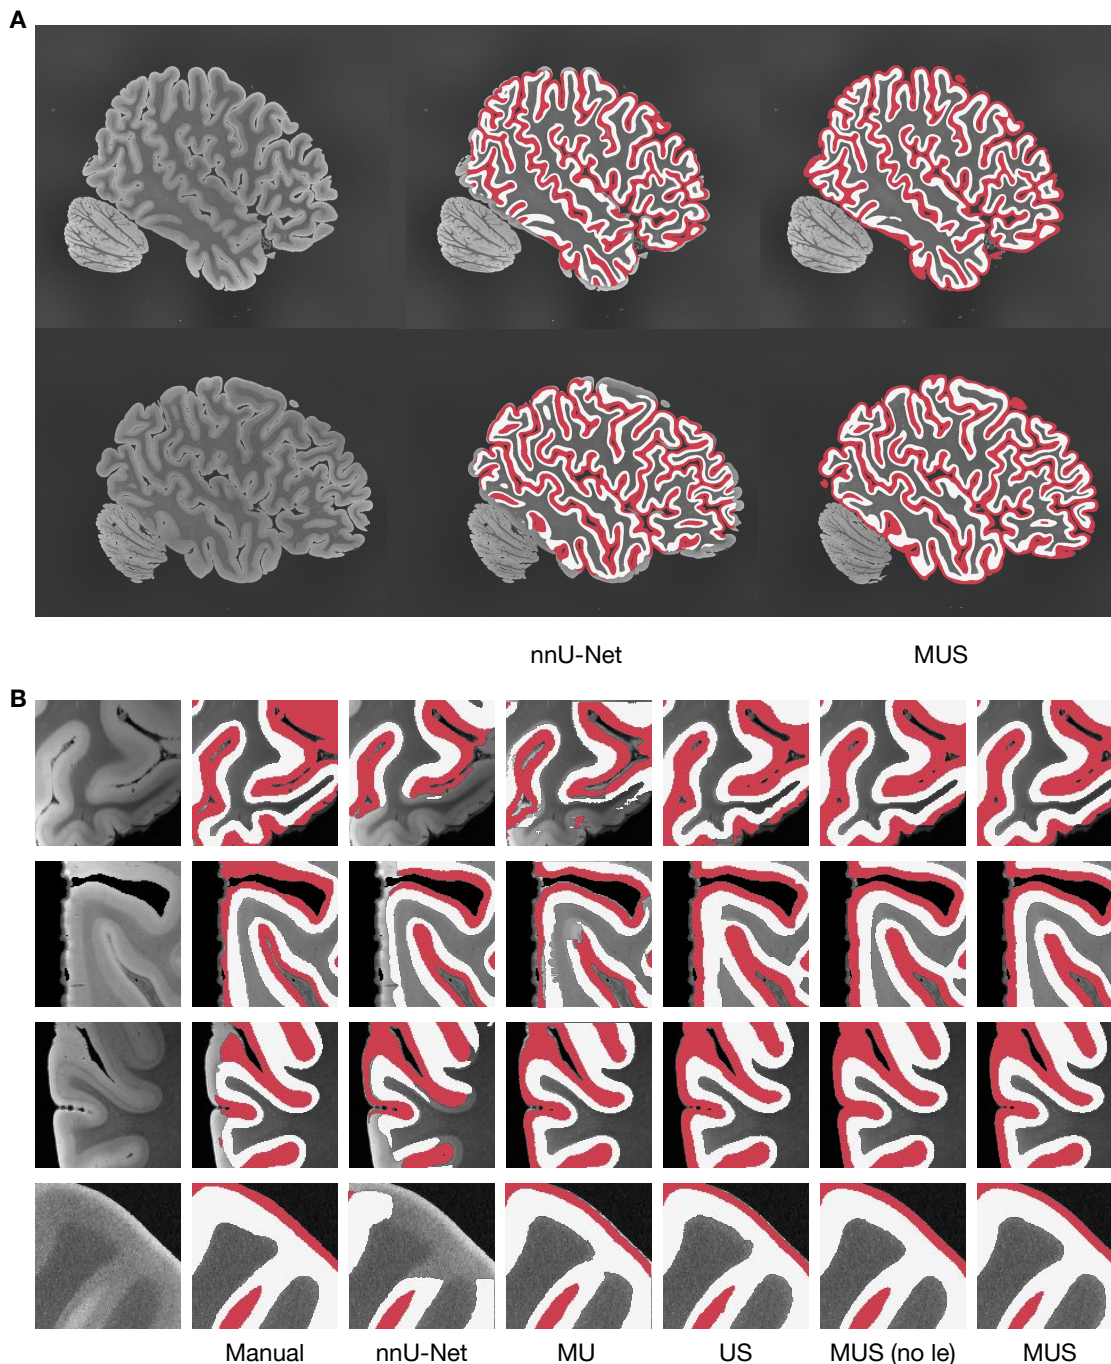
Figure 4: **A.** Example result of layer segmentation (red: supragranular layer, white: infragranular layer) by nnU-Net and our model (MUS: multi-resolution U-Nets semi-supervised). Comprehensive visualization of annotations on all samples can be found in the supplementary document. **B.** Manual annotation and automatic segmentation on example validation slices from our method and baselines (US: simple U-Net semi-supervised; MU: multi-resolution U-Nets supervised; MUS (no l_e): multi-resolution U-Nets semi-supervised with no error diversity loss).

15

3.3 Quantitative segmentation performance as a function of cortical region

We computed the Dice scores for the competing methods and presented them in Table 2. The intra-rater variability was calculated as 0.856 for the supragranular layer and 0.829 for the infragranular layer. Among the methods, MUS demonstrated the highest Dice score, attaining 0.828 for the supragranular layer and 0.818 for the infragranular layer. This latter outcome approaches intra-rater variability, signifying a segmentation performance close to that achieved by human experts through manual segmentation.

Table 2: Dice score result of proposed method and baseline methods, the standard deviation is calculated across samples.

	nnU-Net	MU	US	MUS (no l_e)	MUS	Intra-rater
Supragranular	0.726 ± 0.042	0.783 ± 0.044	0.796 ± 0.051	0.807 ± 0.039	$\mathbf{0.828 \pm 0.040}$	0.856
Infragranular	0.758 ± 0.060	0.769 ± 0.056	0.802 ± 0.048	0.815 ± 0.041	$\mathbf{0.818 \pm 0.037}$	0.829

As expected, the segmentation performance excelled in the BA 44/45 region (Figure 5), which had full labeling in the training dataset. Generally, regions in close anatomical proximity or similar laminar structure to BA 44/45 exhibited good segmentation performance. Regions distant or anatomically dissimilar to BA 44/45, such as the primary visual cortex (V1), entorhinal, and perirhinal cortex, exhibited slightly lower segmentation performance. These findings suggest that incorporating manual segmentations for additional cortical areas may be necessary to enhance the overall supra- and infragranular layer segmentation across the entire hemisphere in future efforts.
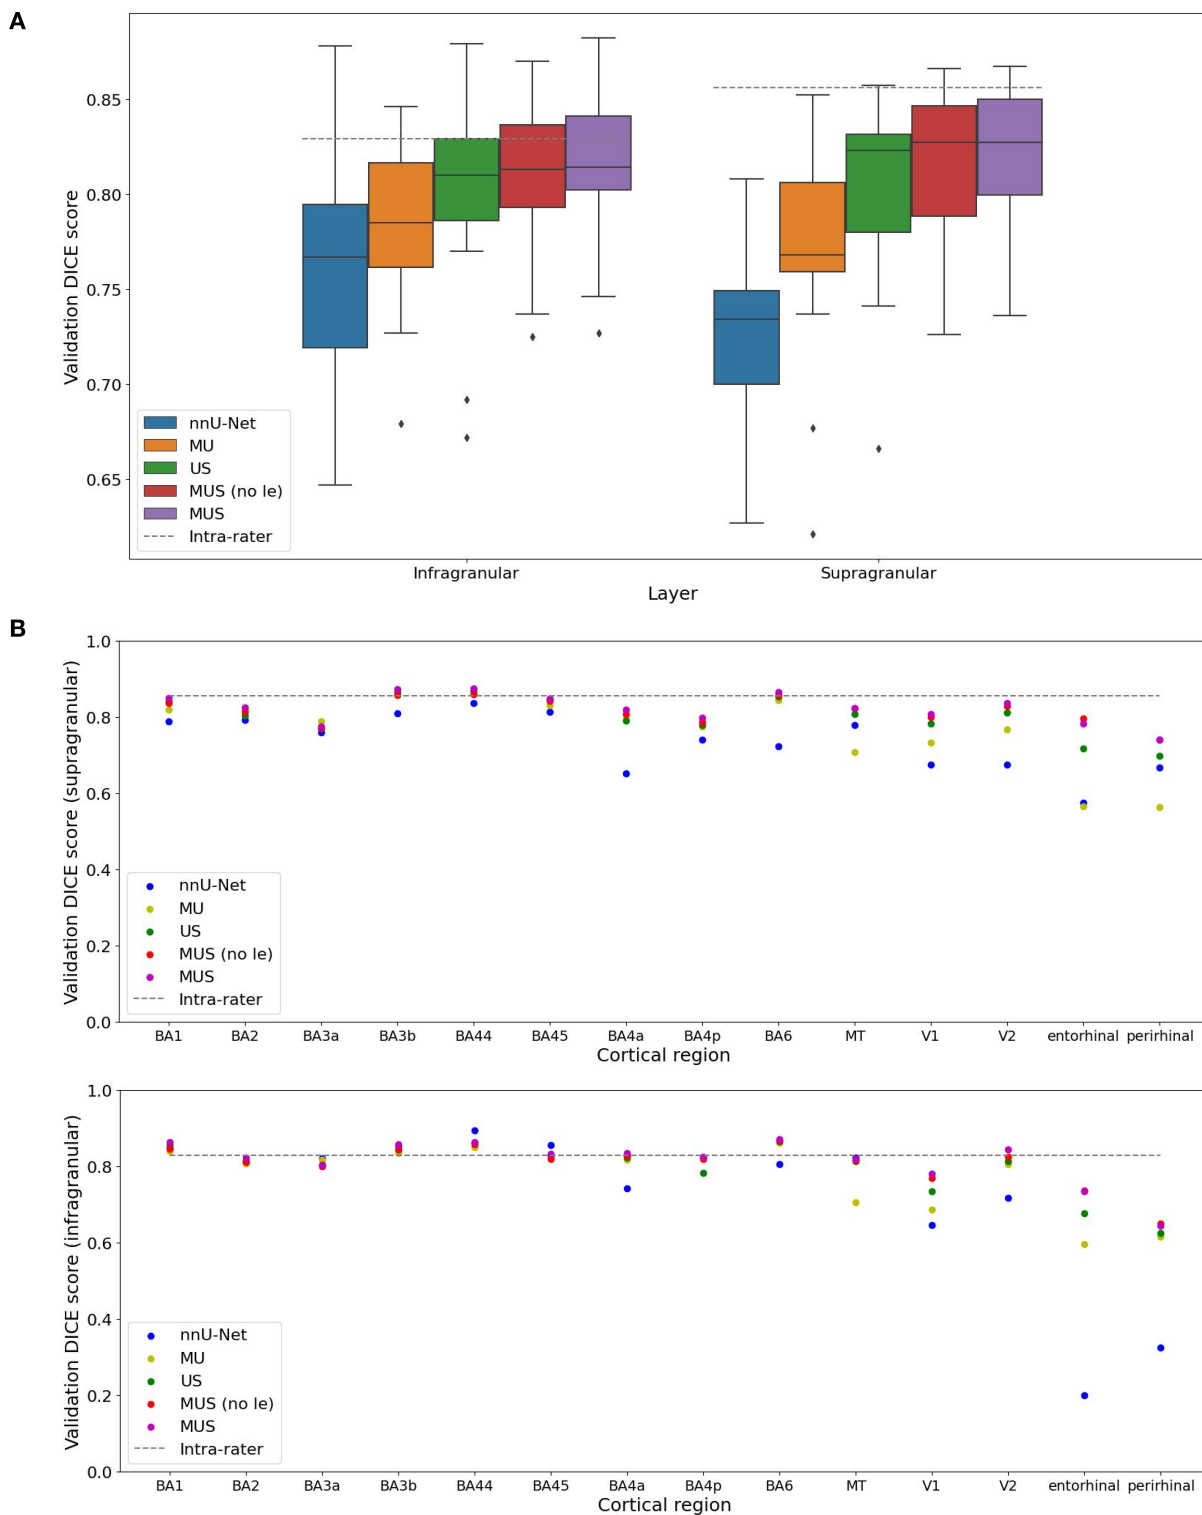
Figure 5: **A.** Box plots of segmentation performance across samples by proposed and baseline methods. **B.** Performance specific to each cortical region (BA3a: somatosensory area (anterior); BA3p: somatosensory area (posterior); MT: visual area (middle temporal); V1: primary visual area; V2: secondary visual area).

3.4   Ablation study

We conducted ablation studies to analyze the importance of three key components in our model design: (1) multi-resolution U-Nets, (2) semi-supervised training, and (3) error diversity loss. In the basic U-Net semi-supervised model, instead of employing a cascade of U-Nets at various resolutions, we utilized a single U-Net operating on patches (size 128^3) at full resolution. In the multi-resolution U-Nets supervised model, we removed the cross pseudo-supervision for the semi-supervised training, ensuring that only labeled training data contributed to the loss calculation. In the multi-resolution U-Nets semi-supervised model without error diversity loss, we excluded the error diversity loss.

Notably, as shown in Table 2, the absence of semi-supervised training led to a reduction in validation Dice score of approximately 0.04. Similarly, without multi-resolution U-Nets, there was an approximate 0.02 decrease in the validation Dice score, and without error diversity loss, the validation Dice score declined by ~ 0.01.

These results indicate that all three components (multi-resolution U-Nets, semi-supervised training, and error diversity loss) contributed to the performance of our segmentation model. Notably, the largest improvement in accuracy was attributed to the implementation of semi-supervised training.

4   Discussion

Accurate segmentation of cortical layers is essential for a comprehensive understanding of neocortical structural organization and its relevance to various neurological conditions and cognitive competencies. The neocortical division into six layers, each characterized by distinct connectivity patterns, underscores the critical importance of precise laminar identification. Leveraging an unprecedented ultra-high-resolution *ex vivo* whole-hemisphere MRI dataset and meticulous but sparse manual annotation, we introduce an innovative approach for the segmentation of supragranular and infragranular layers. For the first time, we obtain a reliable fine segmentation model covering the entire hemisphere.

Our proposed segmentation model, built on an enhanced version of the U-Net architecture and incorporating cross pseudo-supervision, demonstrates remarkable success in accurately delineating supra-and infragranular layers, achieving a Dice score over 0.8. Unlike most existing MRI segmentation models that heavily rely on fully annotated training data and operate at a single resolution, our semi-supervised multi-resolution U-Nets offer a valuable improvement: They reduce the need for large amounts of manually annotated training data and enhance efficiency when processing large volumes in an end-to-end training fashion.

Rigorous ablation studies have demonstrated the efficacy of our novel modules.

Research focusing on supra- and infragranular layers has significant clinical implications. Prior studies have revealed distinct gene expression alterations, pathology accumulations, and atrophies between these layers in patients with conditions such as schizophrenia [4], autism spectrum disorder [61], and epilepsy [62], Alzheimer's disease [63, 64, 65], Parkinson's disease [66], and Huntington's disease [67]. Our high-resolution segmentation maps of these layers across the entire hemisphere will facilitate multiscale investigations of these diseases by integrating with other data types like histological and genomic studies.

The introduced semi-supervised segmentation approach and its corresponding results hold promise for broader applications. It enables benchmarking and validating cortical layer segmentation outcomes across different imaging modalities, fostering cross-modal integration and enriching our understanding of cortical organization. In addition, this method sets the stage for finer segmentation of additional cortical layers and small subcortical nuclei in the future, allowing for even greater granularity in the analysis of cortical and subcortical architecture. Finally, our results can be used to construct surface models, providing insights into alterations in cortical thickness and sulcal depth in diseased states.

While our proposed segmentation model demonstrates promising results, two limitations should be acknowledged. First, the semi-supervised nature of our approach, utilizing a substantial majority of unlabeled data, introduces a degree of uncertainty in the training process. While this approach enhances efficiency, it may also lead to variations in segmentation performance across different cortical regions, as seen in the quantitative analysis. The model's reliance on manual annotations in specific regions may limit its generalizability to areas with sparse or no labeled training data. Second, the proposed model's performance may be influenced by factors such as post-mortem tissue properties, variability in brain morphology, and MRI imaging condition and parameters. Addressing these limitations and conducting further validation on diverse datasets will be crucial for ensuring the robustness and applicability of the presented approach.

In summary, this project presents an advancement in the segmentation of cortical layers within ultra-high-resolution *ex vivo* MRI data. We introduce the first whole-hemisphere segmentation model of sup- and infragranular layers, thereby elevating the delineation of the human cerebral cortex in MRI from a single layer to a dual-layer representation. The incorporation of multi-resolution U-Nets and semi-supervised learning in the segmentation process has demonstrated impressive accuracy and reliability. The potential applications of this segmentation model are extensive, spanning from basic neuroscience research to clinical studies investigating various neurological conditions.

Acknowledgements

Funding

Competing interests

BF has a financial interest in CorticoMetrics, a company whose medical pursuits focus on brain imaging and measurement technologies. BF's interests were reviewed and are managed by Massachusetts General Hospital and Partners HealthCare in accordance with their conflict of interest policies.

References

[1] Yumiko Hatanaka, Yan Zhu, Makio Torigoe, Yoshiaki Kita, and Fujio Murakami. From migration to settlement: the pathways, migration modes and dynamics of neurons in the developing brain. *Proceedings of the Japan Academy, Series B*, 92(1):1–19, 2016.

[2] Xin Tan and Song-Hai Shi. Neocortical neurogenesis and neuronal migration. *Wiley Interdisciplinary Reviews: Developmental Biology*, 2(4):443–459, 2013.

[3] Charles R Gerfen, Michael N Economo, and Jayaram Chandrashekar. Long distance projections of cortical pyramidal neurons. *Journal of neuroscience research*, 96(9):1467–1475, 2018.

[4] Dominique Arion, Szatmár Horváth, David A Lewis, and Károly Mirnics. Infragranular gene expression disturbances in the prefrontal cortex in schizophrenia: signature of altered neural development? *Neurobiology of disease*, 37(3):738–746, 2010.

[5] Ignacio Negrón-Oyarzo, Ariel Lara-Vásquez, Ismael Palacios-García, Pablo Fuentealba, and Francisco Aboitiz. Schizophrenia and reelin: a model based on prenatal stress to study epigenetics, brain development and behavior. *Biological research*, 49(1):1–10, 2016.

[6] Francine M Benes and Sabina Berretta. Gabaergic interneurons: implications for understanding schizophrenia and bipolar disorder. *Neuropsychopharmacology*, 25(1):1–27, 2001.

[7] Jasmin Camacho, Ehsan Ejaz, Jeanelle Ariza, Stephen C Noctor, and Verónica Martínez-Cerdeño. Reln-expressing neuron density in layer I of the superior temporal lobe is similar in human brains with autism and in age-matched controls. *Neuroscience letters*, 579:163–167, 2014.

[8] Rita R Romito-DiGiacomo, Harry Menegay, Samantha A Cicero, and Karl Herrup. Effects of Alzheimer's disease on different cortical layers: the role of intrinsic differences in aβ susceptibility. *Journal of Neuroscience*, 27(32):8496–8504, 2007.

[9] Shan Yang, Zhengyi Yang, Karin Fischer, Kai Zhong, Jörg Stadler, Frank Godenschweger, Johann Steiner, Hans-Jochen Heinze, Hans-Gert Bernstein, Bernhard Bogerts, et al. Integration of ultra-high field MRI and histology for connectome based research of brain disorders. *Frontiers in neuroanatomy*, 7:31, 2013.

[10] Veronica Popescu, Roel Klaver, Adriaan Versteeg, Pieter Voorn, Jos WR Twisk, Frederik Barkhof, Jeroen JG Geurts, and Hugo Vrenken. Postmortem validation of MRI cortical volume measurements in MS. *Human brain mapping*, 37(6):2223–2233, 2016.

[11] Laura E Jonkman, Roel Klaver, Lazar Fleysher, Matilde Inglese, and Jeroen JG Geurts. The substrate of increased cortical fa in ms: A 7T post-mortem MRI and histopathology study. *Multiple Sclerosis Journal*, 22(14):1804–1811, 2016.

[12] Irene Costantini, Leah Morgan, Jiarui Yang, Yael Balbastre, Divya Varadarajan, Luca Pesce, Marina Scardigli, Giacomo Mazzamuto, Vladislav Gavryusev, Filippo Maria Castelli, et al. A cellular resolution atlas of Broca's area. *Science Advances*, 9(41):eadg3844, 2023.

[13] Brian L Edlow, Azma Mareyam, Andreas Horn, Jonathan R Polimeni, Thomas Witzel, M Dylan Tisdall, Jean C Augustinack, Jason P Stockmann, Bram R Diamond, Allison Stevens, et al. 7 Tesla MRI of the ex vivo human brain at 100 micron resolution. *Scientific data*, 6(1):244, 2019.

[14] Max C Keuken, Bethany R Isaacs, Robert Trampel, W Van Der Zwaag, and BU Forstmann. Visualizing the human subcortex using ultra-high field magnetic resonance imaging. *Brain topography*, 31:513–545, 2018.

[15] Kevin R Sitek, Omer Faruk Gulban, Evan Calabrese, G Allan Johnson, Agustin Lage-Castellanos, Michelle Moerel, Satrajit S Ghosh, and Federico De Martino. Mapping the human subcortical auditory system using histology, postmortem MRI and in vivo MRI at 7t. *elife*, 8:e48932, 2019.

[16] Kwok-Shing Chan, Renaud Hédouin, Jeroen Mollink, Jenni Schulz, Anne-Marie van Cappellen van Walsum, and José P Marques. Imaging white matter microstructure with gradient-echo phase imaging: Is ex vivo imaging with formalin-fixed tissue a good approximation of the in vivo brain? *Magnetic Resonance in Medicine*, 88(1):380–390, 2022.

[17] Sanghoon Kim, Ken Sakaie, Ingmar Blümcke, Stephen Jones, and Mark J Lowe. Whole-brain, ultra-high spatial resolution ex vivo MRI with off-the-shelf components. *Magnetic Resonance Imaging*, 76:39–48, 2021.

[18] Pulkit Khandelwal, Michael Tran Duong, Shokufeh Sadaghiani, Sydney Lim, Amanda Denning, Eunice Chung, Sadhana Ravikumar, Sanaz Arezoumandan, Claire Peterson, Madigan Bedard, et al. Automated deep learning segmentation of high-resolution 7 T

ex vivo MRI for quantitative analysis of structure-pathology correlations in neurodegenerative diseases. *arXiv preprint arXiv:2303.12237*, 2023.

[19] Katrin Amunts, Claude Lepage, Louis Borgeat, Hartmut Mohlberg, Timo Dickscheid, Marc-Étienne Rousseau, Sebastian Bludau, Pierre-Louis Bazin, Lindsay B Lewis, Ana-Maria Oros-Peusquens, et al. Bigbrain: an ultrahigh-resolution 3D human brain model. *science*, 340(6139):1472–1475, 2013.

[20] Song-Lin Ding, Joshua J Royall, Susan M Sunkin, Lydia Ng, Benjamin AC Facer, Phil Lesnar, Angie Guillozet-Bongaarts, Bergen McMurray, Aaron Szafer, Tim A Dolbeare, et al. Comprehensive cellular-resolution atlas of the adult human brain. *Journal of Comparative Neurology*, 524(16):3127–3481, 2016.

[21] Konrad Wagstyl, Stéphanie Larocque, Guillem Cucurull, Claude Lepage, Joseph Paul Cohen, Sebastian Bludau, Nicola Palomero-Gallagher, Lindsay B Lewis, Thomas Funck, Hannah Spitzer, et al. Bigbrain 3D atlas of cortical layers: Cortical and laminar thickness gradients diverge in sensory and motor cortices. *PLoS biology*, 18(4):e3000678, 2020.

[22] Juan Eugenio Iglesias, Ricardo Insausti, Garikoitz Lerma-Usabiaga, Martina Bocchetta, Koen Van Leemput, Douglas N Greve, Andre Van der Kouwe, Bruce Fischl, César Caballero-Gaudes, Pedro M Paz-Alonso, et al. A probabilistic atlas of the human thalamic nuclei combining ex vivo MRI and histology. *Neuroimage*, 183:314–326, 2018.

[23] Zeynep M Saygin, Dorit Kliemann, Juan Eugenio Iglesias, André JW van der Kouwe, Ellen Boyd, Martin Reuter, Allison Stevens, Koenraad Van Leemput, Alexis McKee, Matthew P Frosch, et al. High-resolution magnetic resonance imaging reveals nuclei of the human amygdala: manual segmentation to automatic atlas. *Neuroimage*, 155:370–382, 2017.

[24] Xuxin Chen, Ximin Wang, Ke Zhang, Kar-Ming Fung, Theresa C Thai, Kathleen Moore, Robert S Mannel, Hong Liu, Bin Zheng, and Yuchen Qiu. Recent advances and clinical applications of deep learning in medical image analysis. *Medical Image Analysis*, 79:102444, 2022.

[25] Yichi Zhang, Qingcheng Liao, Le Ding, and Jicong Zhang. Bridging 2D and 3D segmentation networks for computation-efficient volumetric medical image segmentation: An empirical study of 2.5 D solutions. *Computerized Medical Imaging and Graphics*, 99:102088, 2022.

[26] Konstantinos Kamnitsas, Christian Ledig, Virginia FJ Newcombe, Joanna P Simpson, Andrew D Kane, David K Menon, Daniel Rueckert, and Ben Glocker. Efficient multi-scale 3D cnn with fully connected CRF for accurate brain lesion segmentation. *Medical image analysis*, 36:61–78, 2017.

[27] Christian S Perone, Evan Calabrese, and Julien Cohen-Adad. Spinal cord gray matter segmentation using deep dilated convolutions. *Scientific reports*, 8(1):5966, 2018.

[28] Robin Brügger, Christian F Baumgartner, and Ender Konukoglu. A partially reversible U-Net for memory-efficient volumetric image segmentation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part III 22*, pages 429–437. Springer, 2019.

[29] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnU-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021.

[30] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.

[31] Bruce Fischl, David H Salat, André JW Van Der Kouwe, Nikos Makris, Florent Ségonne, Brian T Quinn, and Anders M Dale. Sequence-independent segmentation of magnetic resonance images. *Neuroimage*, 23:S69–S84, 2004.

[32] Divya Varadarajan, Robert Frost, Andre van der Kouwe, Leah Morgan, Bram Diamond, Emma Boyd, Morgan Fogarty, Allison Stevens, Bruce Fischl, and Jonathan R Polimeni. Edge-preserving B0 inhomogeneity distortion correction for high-resolution multi-echo ex vivo MRI at 7T. In *International Society for Magnetic Resonance in Medicine*, volume 664, 2020.

[33] Divya Varadarajan, Mukund Balasubramanian, Daniel J Park, Thomas Witzel, Jason P Stockmann, and Jonathan R Polimeni. Characterizing the acquisition protocol dependencies of B0 field mapping and the effects of eddy currents and spoiling. In *Proc. Int. Soc. Magn. Reson. Med*, volume 29, page 3552, 2021.

[34] Oula Puonti, Juan Eugenio Iglesias, and Koen Van Leemput. Fast and sequence-adaptive whole-brain segmentation using parametric bayesian modeling. *NeuroImage*, 143:235–249, 2016.

[35] Luca Pesce, Marina Scardigli, Vladislav Gavryusev, Annunziatina Laurino, Giacomo Mazzamuto, Niamh Brady, Giuseppe Sancataldo, Ludovico Silvestri, Christophe Destrieux, Patrick R Hof, et al. 3D molecular phenotyping of cleared human brain tissues with light-sheet fluorescence microscopy. *Communications Biology*, 5(1):447, 2022.

[36] Shuaibin Chang, Divya Varadarajan, Jiarui Yang, Ichun Anderson Chen, Sreekanth Kura, Caroline Magnain, Jean C Augustinack, Bruce Fischl, Douglas N Greve, David A Boas, et al. Scalable mapping of myelin and neuron density in the human brain with micrometer resolution. *Scientific Reports*, 12(1):363, 2022.

[37] Matthew F Glasser, Manu S Goyal, Todd M Preuss, Marcus E Raichle, and David C Van Essen. Trends and properties of human cerebral cortex: correlations with cortical myelin content. *Neuroimage*, 93:165–175, 2014.

[38] Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.

[39] Benjamin Billot, Douglas N Greve, Oula Puonti, Axel Thielscher, Koen Van Leemput, Bruce Fischl, Adrian V Dalca, Juan Eugenio Iglesias, et al. Synthseg: Segmentation of brain MRI scans of any contrast and resolution without retraining. *Medical image analysis*, 86:102789, 2023.

[40] Pavlo Radiuk. Applying 3D U-net architecture to the task of multi-organ segmentation in computed tomography. 2020.

[41] Alexander Ushinsky, Michelle Bardis, Justin Glavis-Bloom, Edward Uchio, Chanon Chantaduly, Michael Nguyentat, Daniel Chow, Peter D Chang, and Roozbeh Houshyar. A 3D-2D hybrid U-net convolutional neural network approach to prostate organ segmentation of multiparametric MRI. *American journal of roentgenology*, 216(1):111–116, 2021.

[42] Nagwa M Aboelenein, Piao Songhao, Anis Koubaa, Alam Noor, and Ahmed Afifi. HTTU-Net: Hybrid two track U-Net for automatic brain tumor segmentation. *IEEE Access*, 8:101406–101415, 2020.

[43] Maayan Frid-Adar, Avi Ben-Cohen, Rula Amer, and Hayit Greenspan. Improving the segmentation of anatomical structures in chest radiographs using U-Net with an imagenet pre-trained encoder. In *Image Analysis for Moving Organ, Breast, and Thoracic*

Images: Third International Workshop, RAMBO 2018, Fourth International Workshop, BIA 2018, and First International Workshop, TIA 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16 and 20, 2018, Proceedings 3, pages 159–168. Springer, 2018.

[44] Abhijit Guha Roy, Sailesh Conjeti, Nassir Navab, Christian Wachinger, Alzheimer's Disease Neuroimaging Initiative, et al. Quicknat: A fully convolutional network for quick and accurate segmentation of neuroanatomy. *NeuroImage*, 186:713–727, 2019.

[45] Fan Jia, Wing Hong Wong, and Tieyong Zeng. DDUNet: Dense dense U-Net with applications in image denoising. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 354–364, 2021.

[46] Mattias P Heinrich, Maik Stille, and Thorsten M Buzug. Residual U-Net convolutional neural network architecture for low-dose ct denoising. *Current Directions in Biomedical Engineering*, 4(1):297–300, 2018.

[47] Zhangpei Cheng, Kaixuan Guo, Changfeng Wu, Jiankun Shen, and Lei Qu. U-Net cascaded with dilated convolution for medical image registration. In *2019 Chinese Automation Congress (CAC)*, pages 3647–3651. IEEE, 2019.

[48] Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019.

[49] Ning Han, Li Zhou, Zhengmao Xie, Jingli Zheng, and Liuxin Zhang. Multi-level U-Net network for image super-resolution reconstruction. *Displays*, 73:102192, 2022.

[50] Zhengyang Lu and Ying Chen. Single image super-resolution based on a modified U-Net with mixed gradient loss. *signal, image and video processing*, pages 1–9, 2022.

[51] Juan E Iglesias, Benjamin Billot, Yaël Balbastre, Colin Magdamo, Steven E Arnold, Sudeshna Das, Brian L Edlow, Daniel C Alexander, Polina Golland, and Bruce Fischl. Synthsr: A public AI tool to turn heterogeneous clinical brain scans into high-resolution T1-weighted images for 3D morphometry. *Science advances*, 9(5):eadd3607, 2023.

[52] Konstantinos Kamnitsas, Enzo Ferrante, Sarah Parisot, Christian Ledig, Aditya V Nori, Antonio Criminisi, Daniel Rueckert, and Ben Glocker. Deepmedic for brain tumor segmentation. In *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: Second International Workshop, BrainLes 2016, with the Challenges on BRATS,*

ISLES and mTOP 2016, Held in Conjunction with MICCAI 2016, Athens, Greece, October 17, 2016, Revised Selected Papers 2, pages 138–149. Springer, 2016.

[53] Han Zheng, Lanfen Lin, Hongjie Hu, Qiaowei Zhang, Qingqing Chen, Yutaro Iwamoto, Xianhua Han, Yen-Wei Chen, Ruofeng Tong, and Jian Wu. Semi-supervised segmentation of liver using adversarial learning with deep atlas prior. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part VI 22*, pages 148–156. Springer, 2019.

[54] Sukesh Adiga Vasudeva, Jose Dolz, and Herve Lombaert. Leveraging labeling representations in uncertainty-based semi-supervised segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 265–275. Springer, 2022.

[55] Yassine Ouali, Céline Hudelot, and Myriam Tami. Semi-supervised semantic segmentation with cross-consistency training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12674–12684, 2020.

[56] Gerda Bortsova, Florian Dubost, Laurens Hogeweg, Ioannis Katramados, and Marleen De Bruijne. Semi-supervised medical image segmentation via learning consistency under transformations. In *Medical Image Computing and Computer Assisted Intervention– MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part VI 22*, pages 810–818. Springer, 2019.

[57] Xiaokang Chen, Yuhui Yuan, Gang Zeng, and Jingdong Wang. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2613–2622, 2021.

[58] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. Ieee, 2016.

[59] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

[60] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[61] Aaron T Karst and Jeffrey J Hutsler. Two-dimensional analysis of the supragranular layers in autism spectrum disorder. *Research in Autism Spectrum Disorders*, 32:96–105, 2016.

[62] Kinga Tóth, Katharina T Hofer, Ágnes Kandrács, László Entz, Attila Bagó, Loránd Erőss, Zsófia Jordán, Gábor Nagy, András Sólyom, Dániel Fabó, et al. Hyperexcitability of the network contributes to synchronization processes in the human epileptic neocortex. *The Journal of physiology*, 596(2):317–342, 2018.

[63] Patrick R Hof, N Archin, AP Osmand, JH Dougherty, C Wells, C Bouras, and JH Morrison. Posterior cortical atrophy in Alzheimer's disease: analysis of a new case and re-evaluation of a historical report. *Acta neuropathologica*, 86(3):215–223, 1993.

[64] Patrick R Hof, John H Morrison, and Kevin Cox. Quantitative analysis of a vulnerable subset of pyramidal neurons in Alzheimer's disease: I. superior frontal and inferior temporal cortex. *Journal of Comparative Neurology*, 301(1):44–54, 1990.

[65] Patrick R Hof and John H Morrison. Quantitative analysis of a vulnerable subset of pyramidal neurons in Alzheimer's disease: II. primary and secondary visual cortex. *Journal of Comparative Neurology*, 301(1):55–64, 1990.

[66] YY Fathy, AJ Jonker, E Oudejans, FJJ de Jong, A-MW van Dam, AJM Rozemuller, and WDJ van de Berg. Differential insular cortex subregional vulnerability to α-synuclein pathology in Parkinson's disease and dementia with lewy bodies. *Neuropathology and applied neurobiology*, 45(3):262–277, 2019.

[67] Helmut Heinsen, M Strik, M Bauer, K Luther, G Ulmar, D Gangnus, G Jungkunz, W Eisenmengers, and M Götz. Cortical and striatal neurone number in Huntington's disease. *Acta neuropathologica*, 88(4):320–333, 1994.

Supplementary: Segmentation of supragranular and infragranular layers in ultra-high resolution 7T *ex vivo* MRI of the human cerebral cortex

Xiangrui Zeng[1,2], Oula Puonti[1,3], Areej Sayeed[1,2], Rogeny Herisse[1,2], Jocelyn Mora[1,2], Kathryn Evancic[1,2], Divya Varadarajan[1,2], Yael Balbastre[1,2], Irene Costantini[4,5,6], Marina Scardigli[5], Josephine Ramazzotti[5], Danila DiMeo[5], Giacomo Mazzamuto[4,5,7], Luca Pesce[5], Niamh Brady[5], Franco Cheli[5], Francesco Saverio Pavone[4,5,7], Patrick R. Hof[8], Robert Frost[1,2], Jean Augustinack[1,2], André van der Kouwe[1,2], Juan Eugenio Iglesias[1,2,*], and Bruce Fischl[1,2,*]

[1]Department of Radiology, Athinoula A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Charlestown, MA, USA
[2]Harvard Medical School, Department of Radiology, Boston, MA, USA
[3]Danish Research Centre for Magnetic Resonance, Centre for Functional and Diagnostic Imaging and Research, Copenhagen University Hospital - Amager and Hvidovre, Copenhagen, Denmark
[4]National Research Council - National Institute of Optics (CNR-INO), Sesto Fiorentino, Italy
[5]European Laboratory for Non-Linear Spectroscopy (LENS), Sesto Fiorentino, Italy
[6]Department of Biology, University of Florence, Italy
[7]Department of Physics and Astronomy, University of Florence, Italy
[8]Nash Family Department of Neuroscience and Friedman Brain Institute, Icahn School of Medicine at Mount Sinai, New York, NY, USA
[9]Centre for Medical Image Computing, University College London, London, UK
[10]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, USA
*Corresponding authors. Email address: jiglesiasgonzalez@mgh.harvard.edu, fischl@nmr.mgh.harvard.edu

1

Figure 1: 2D slices of case 1 in sagittal, coronal, and axial views. First row: MRI scan; Second row: manual annotation (red: supragranular layer, white: infragranular layer) overlaid; Third row: MUS prediction overlaid.
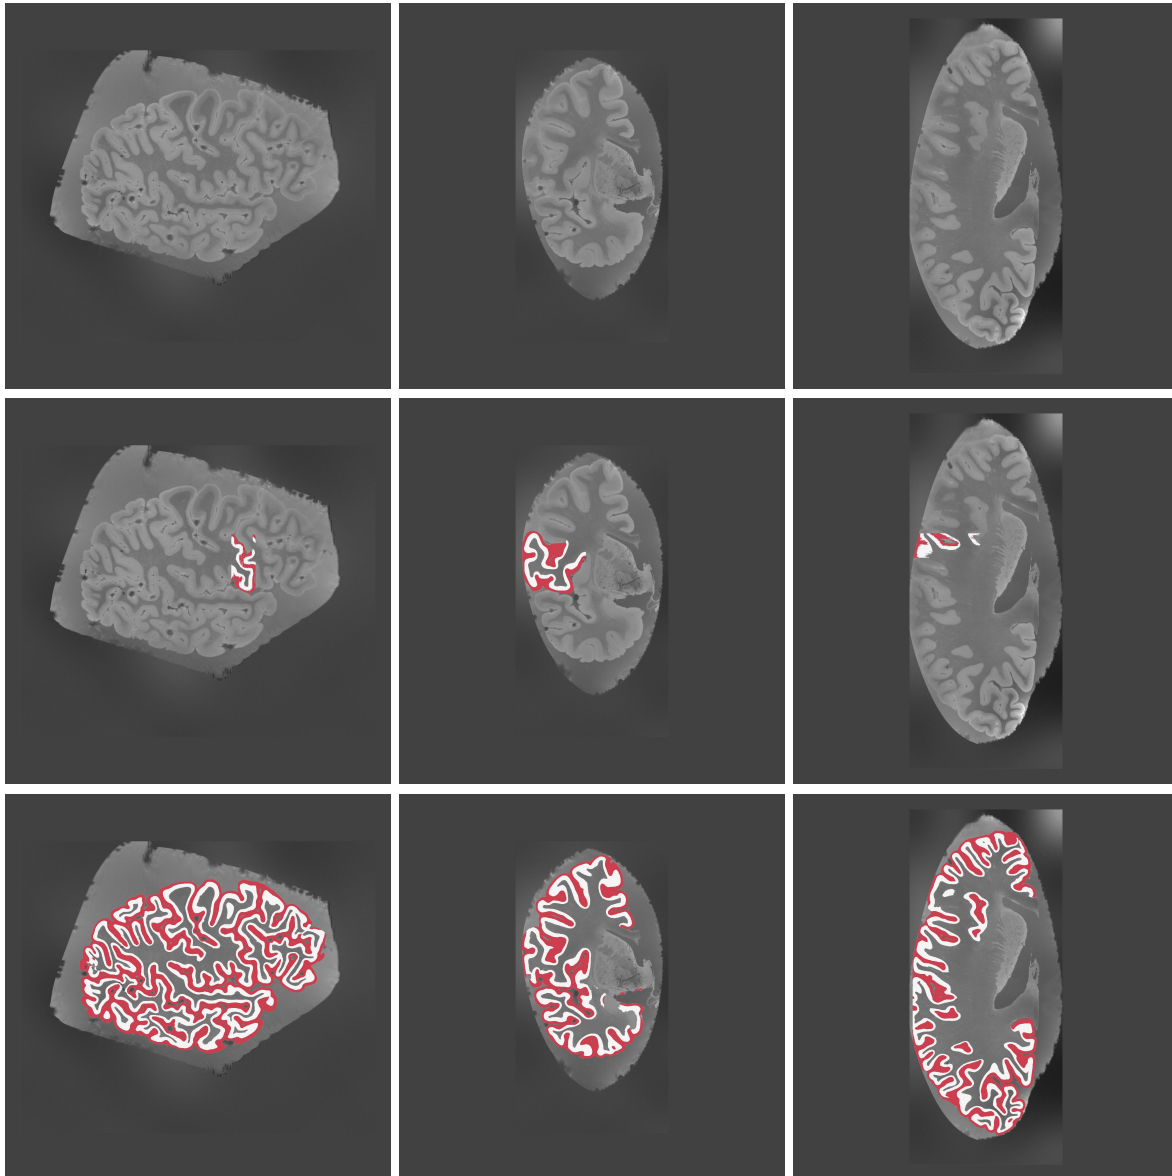
Figure 2: 2D slices of case 2 in sagittal, coronal, and axial views.
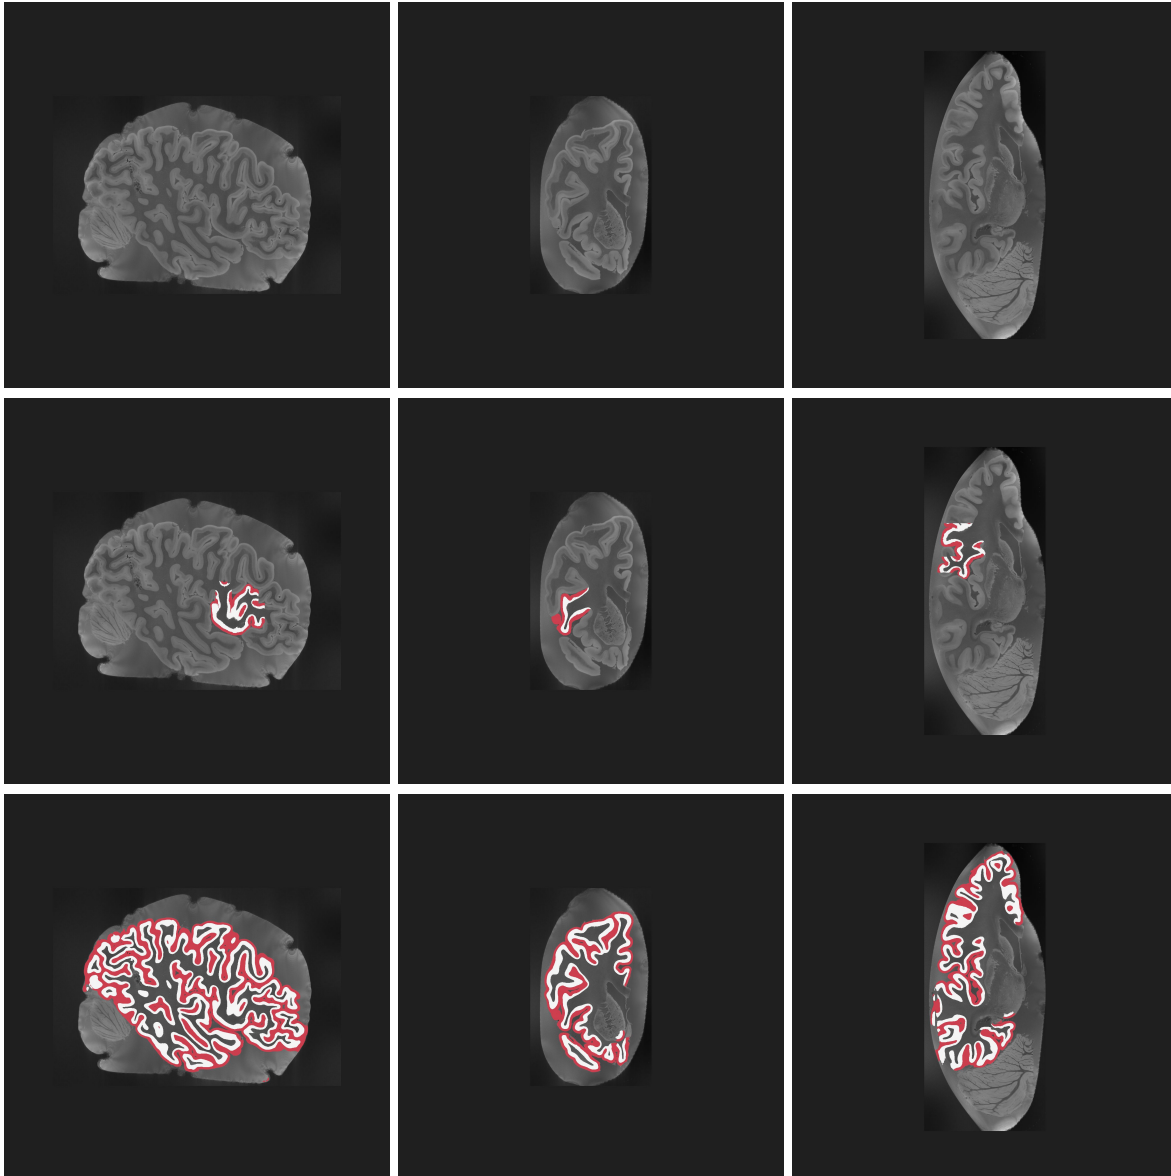
Figure 3: 2D slices of case 3 in sagittal, coronal, and axial views.
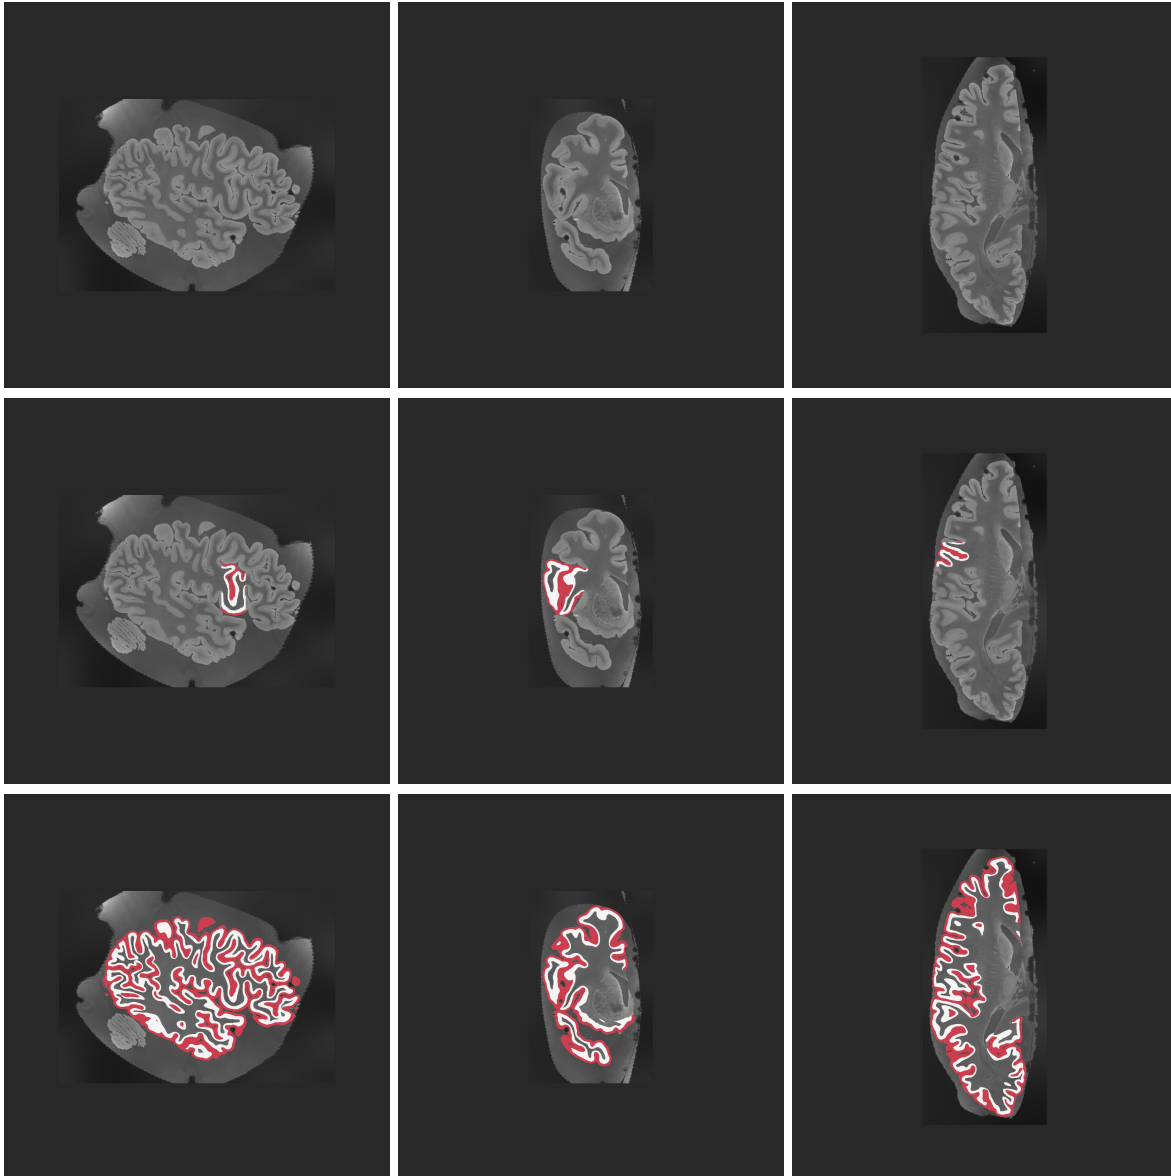
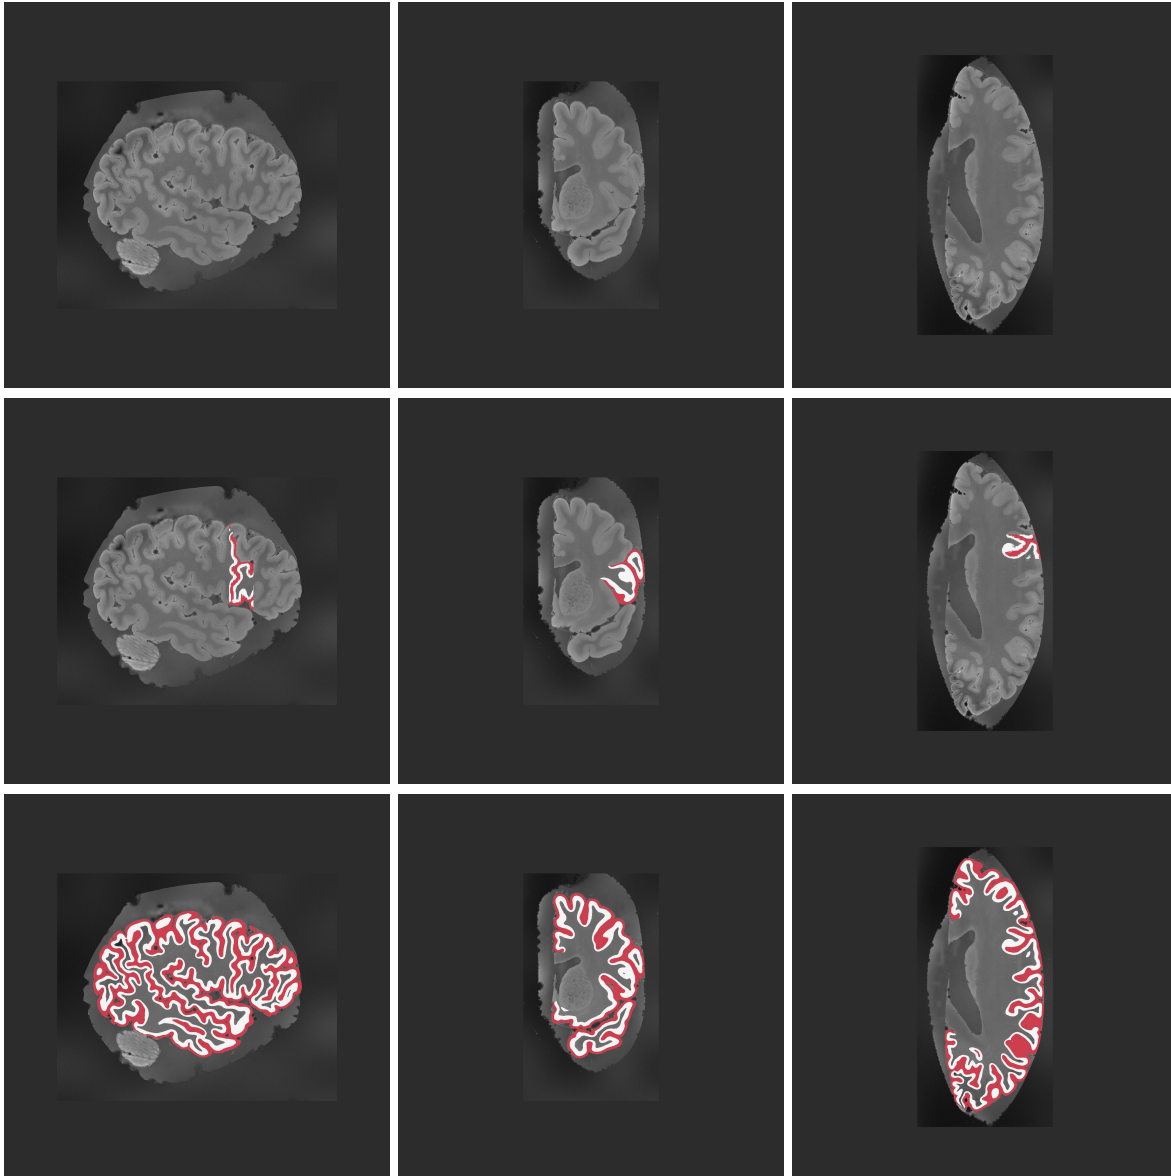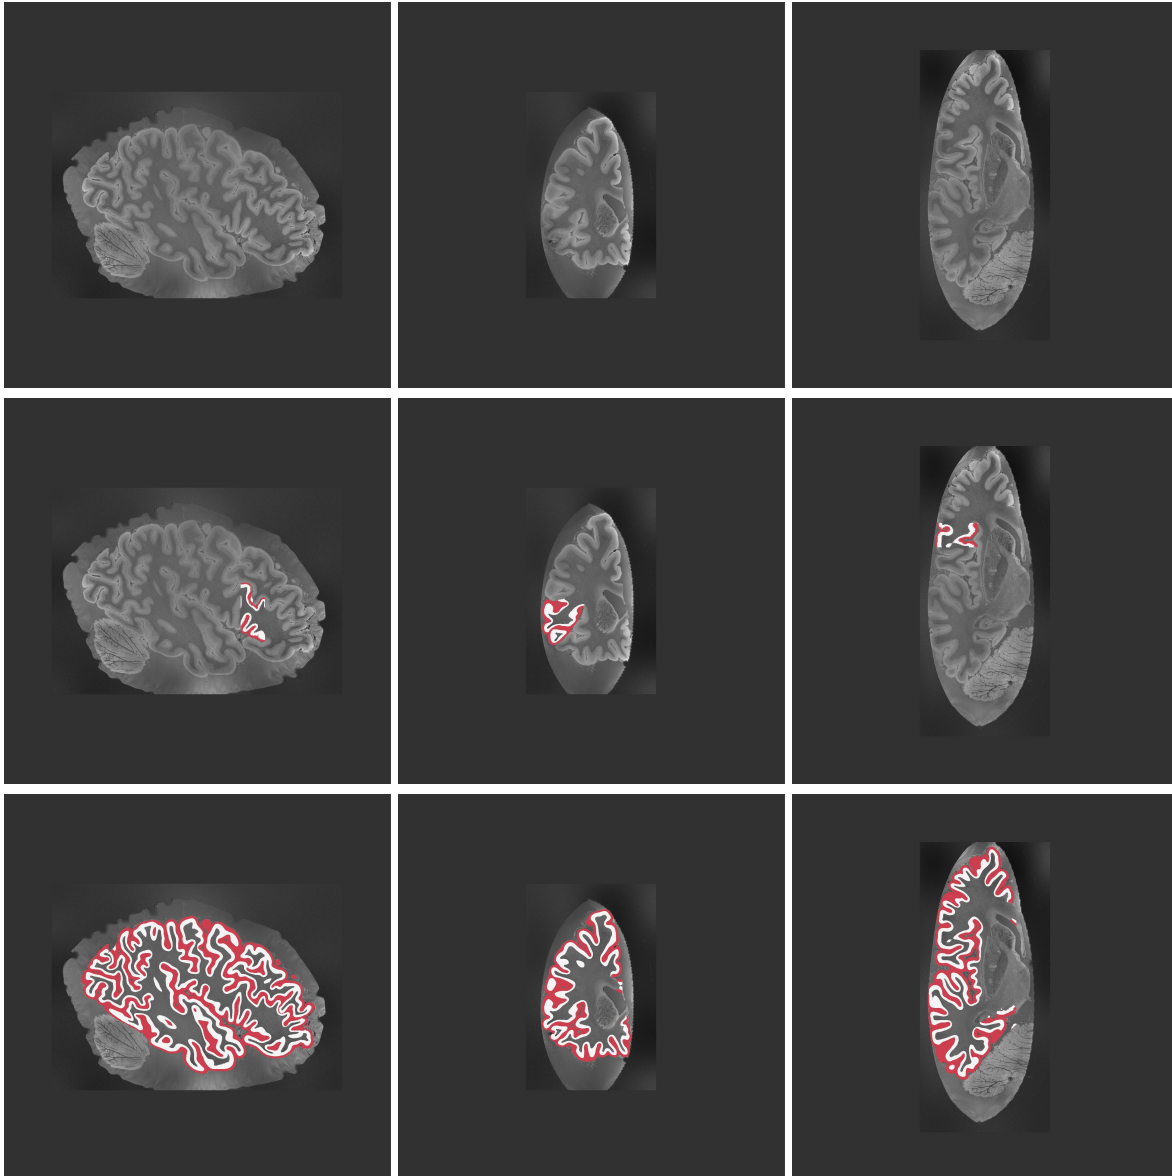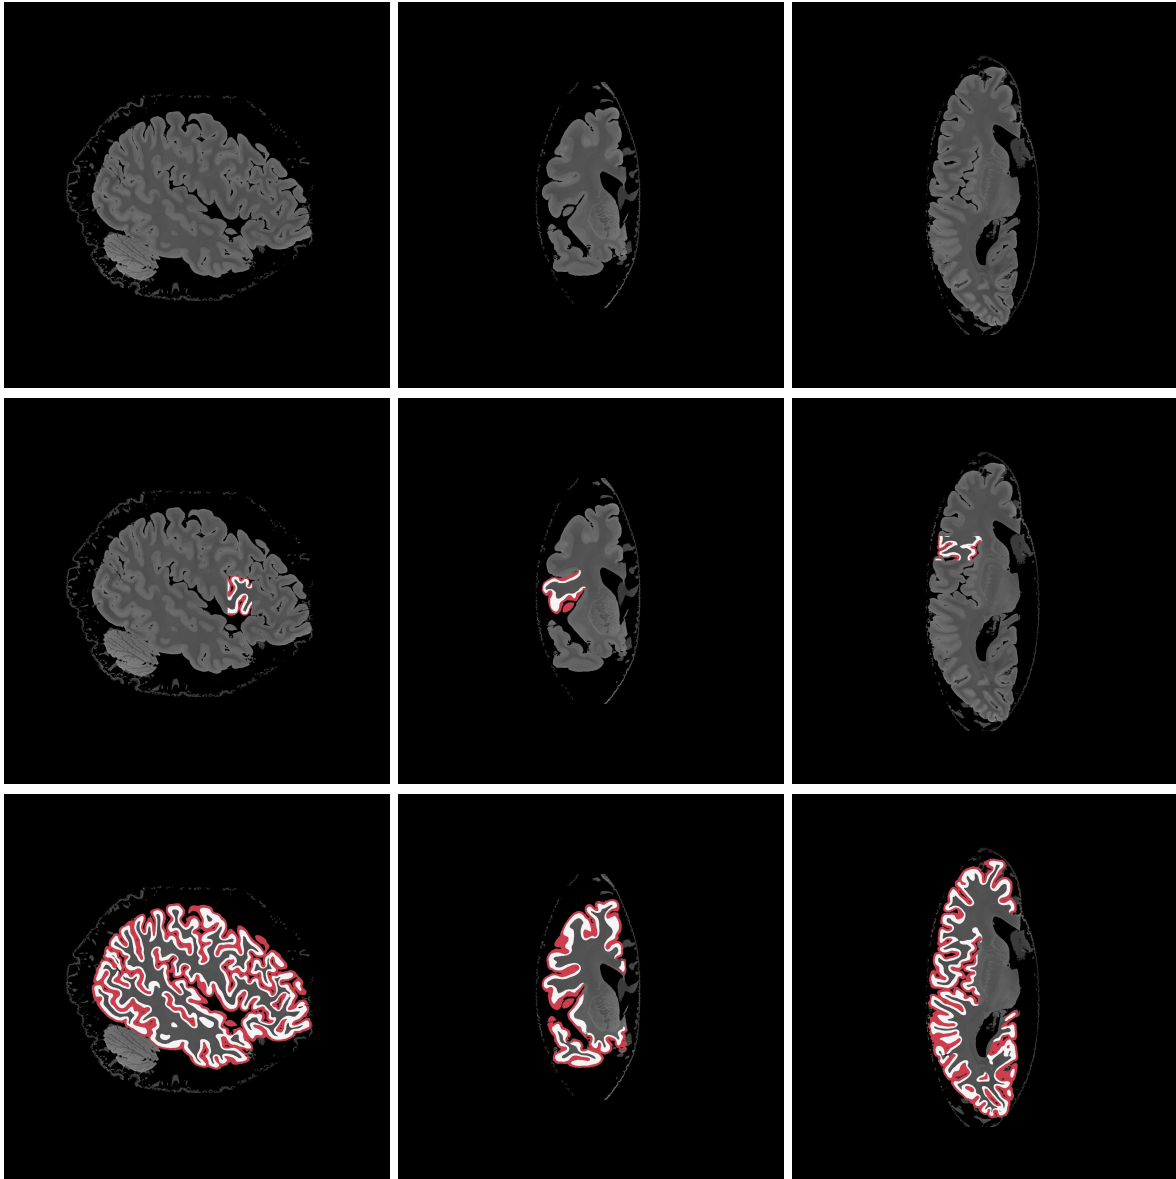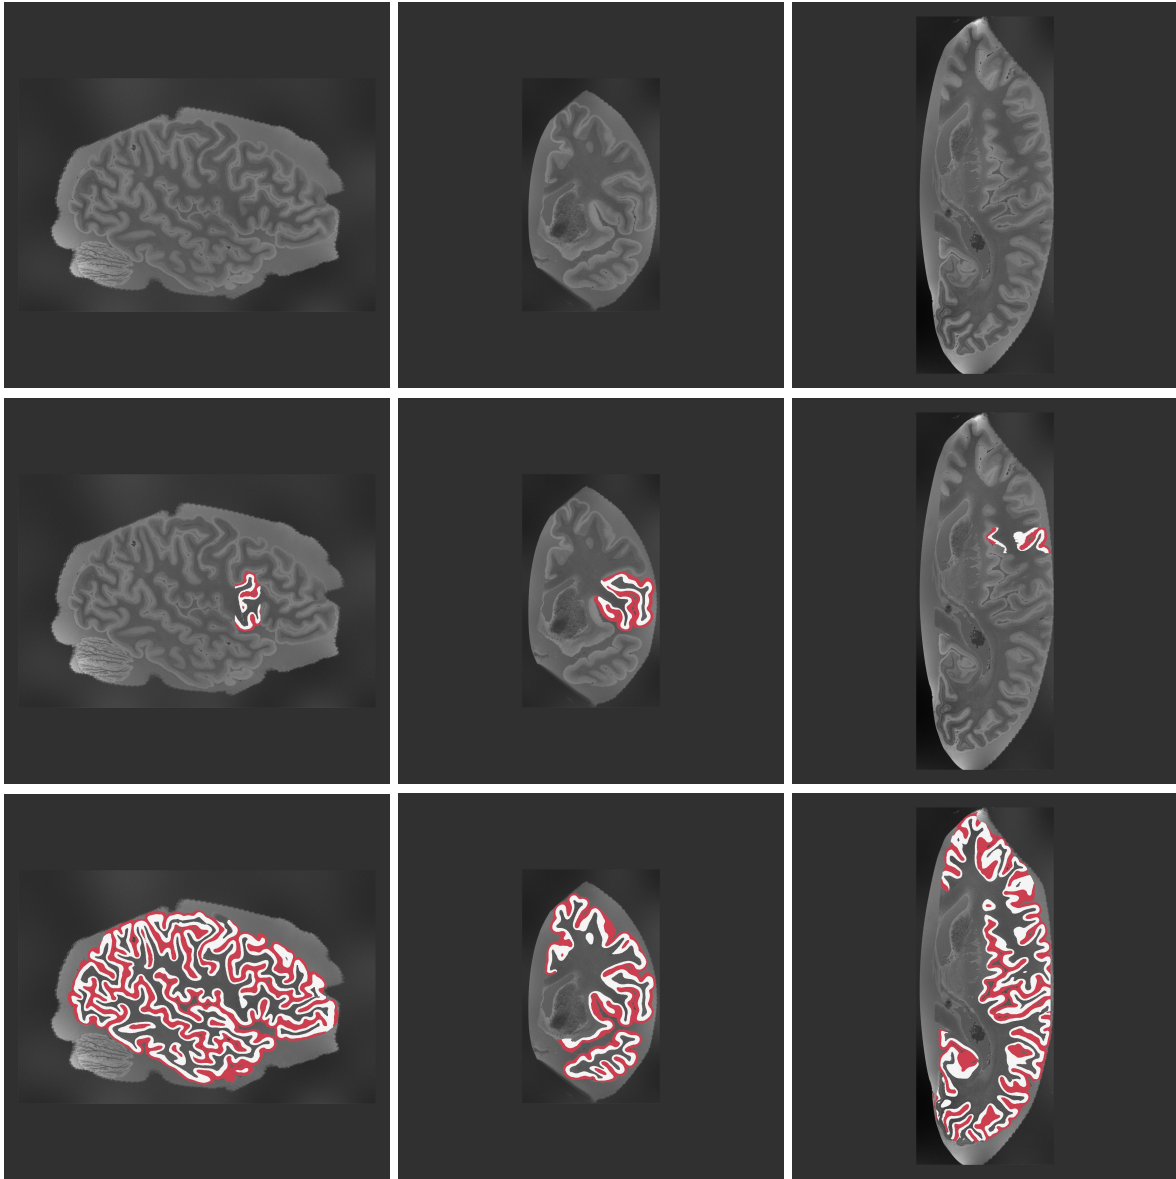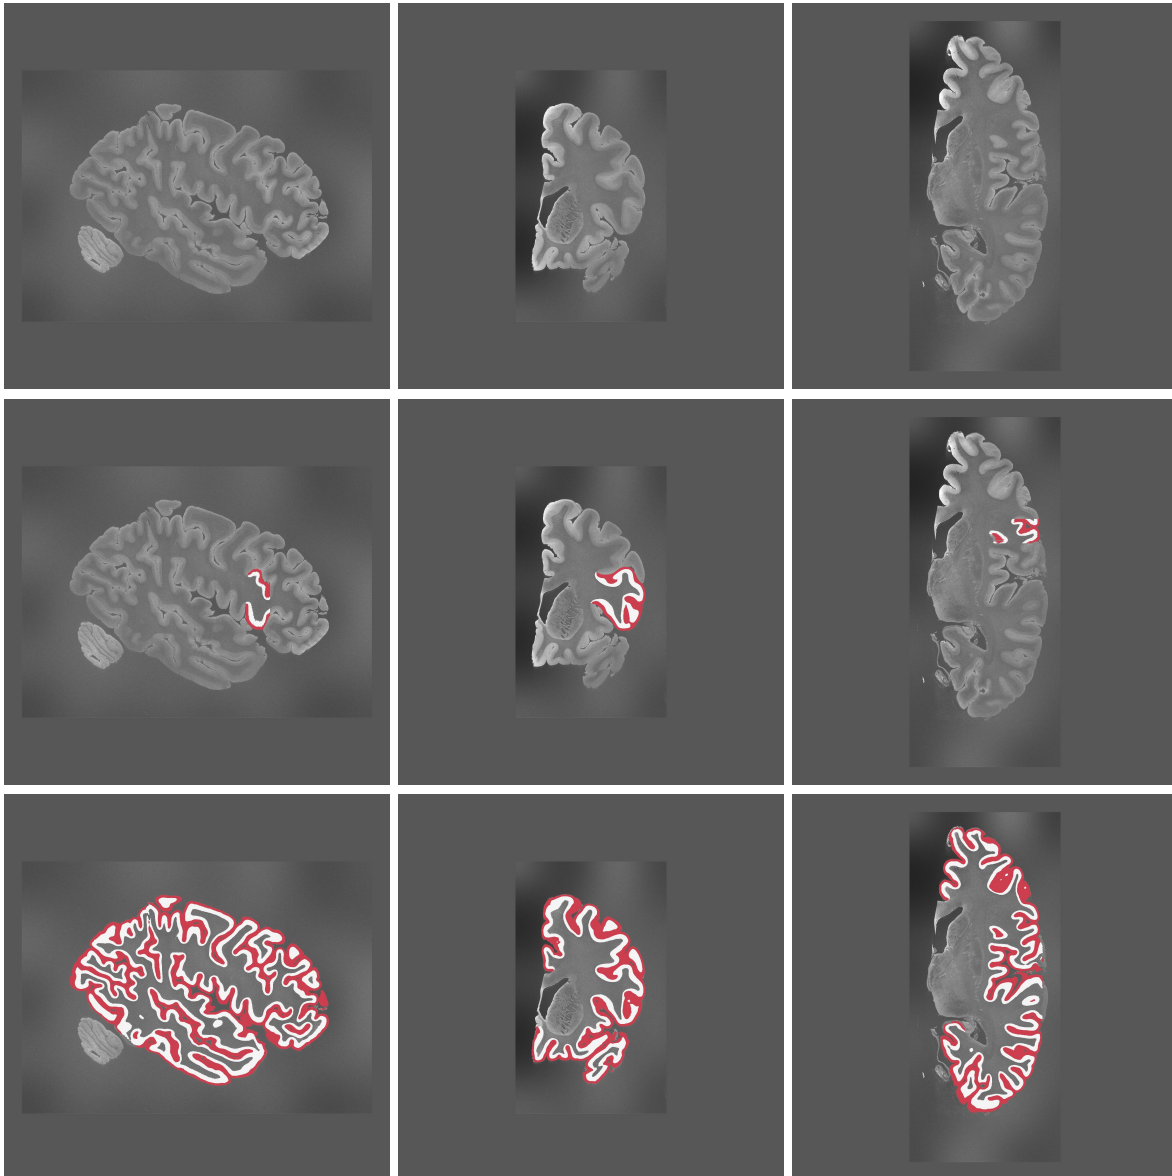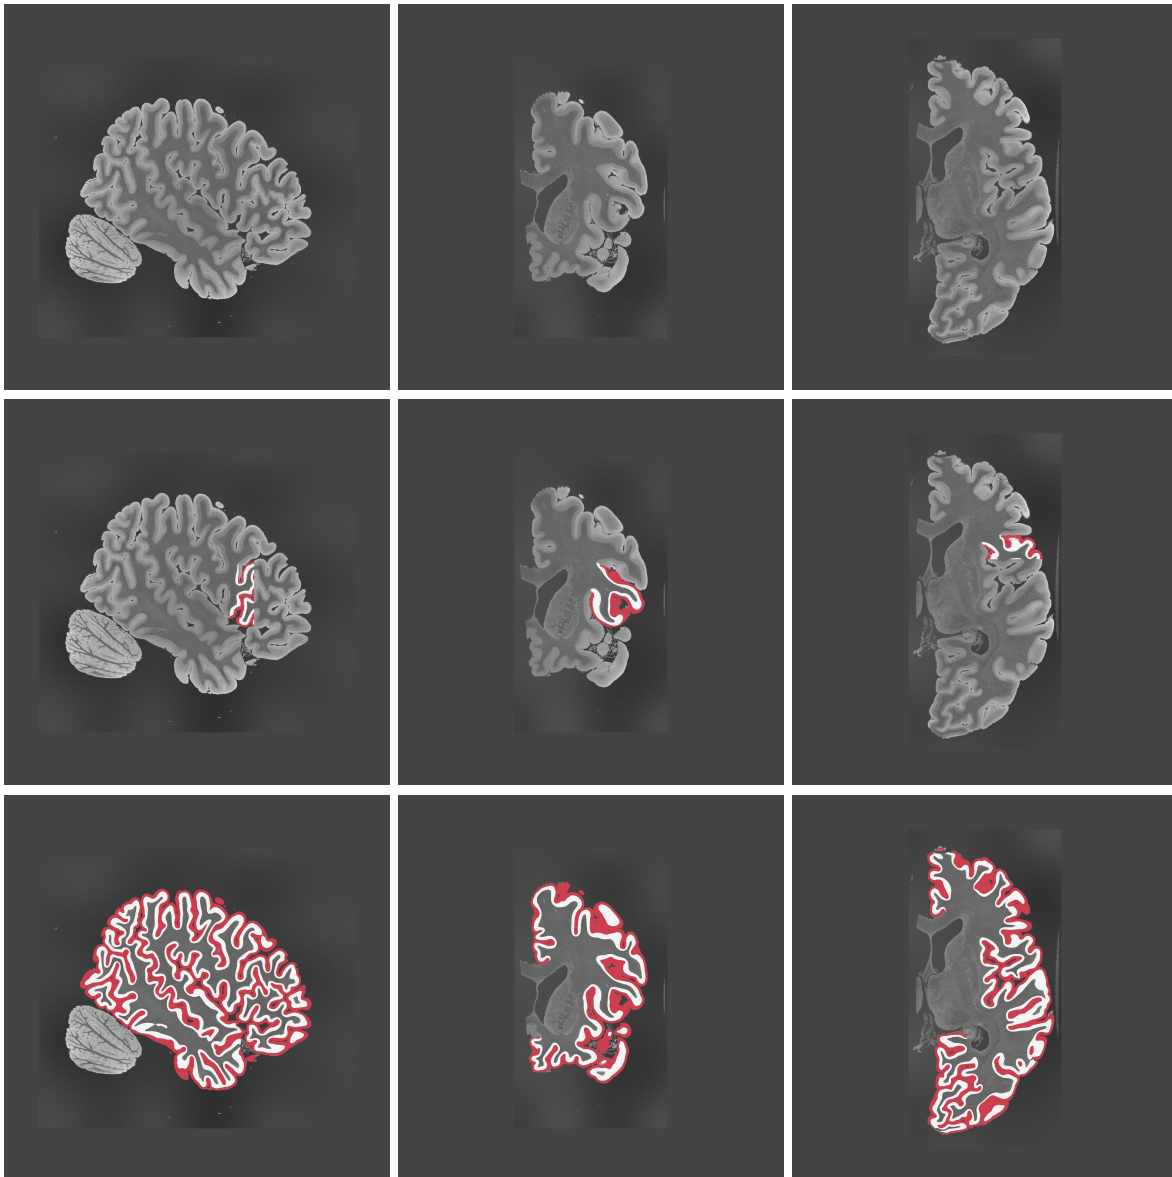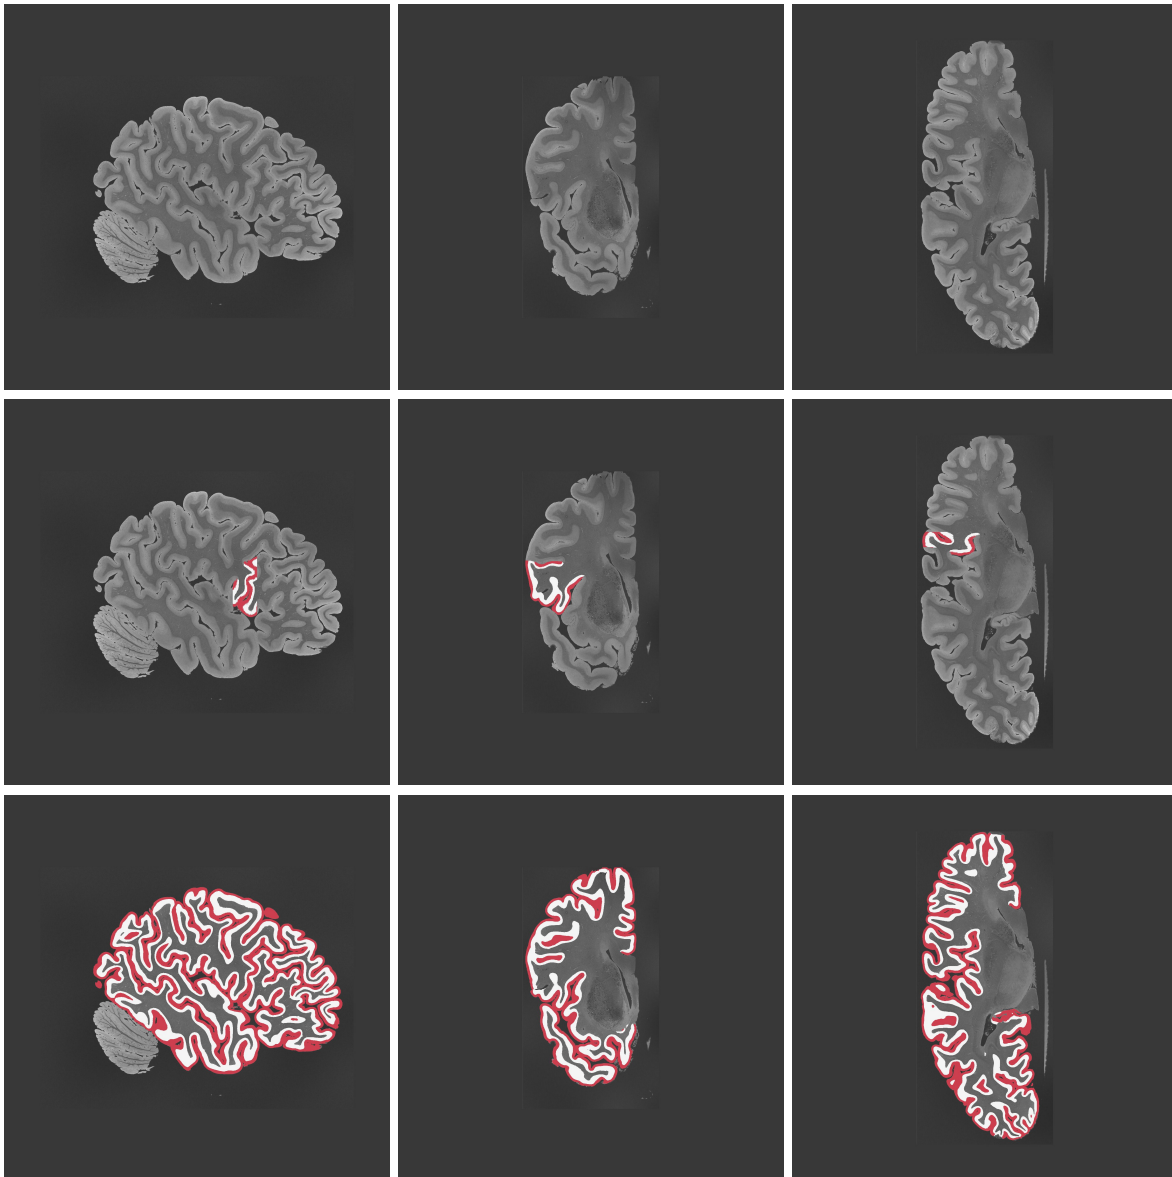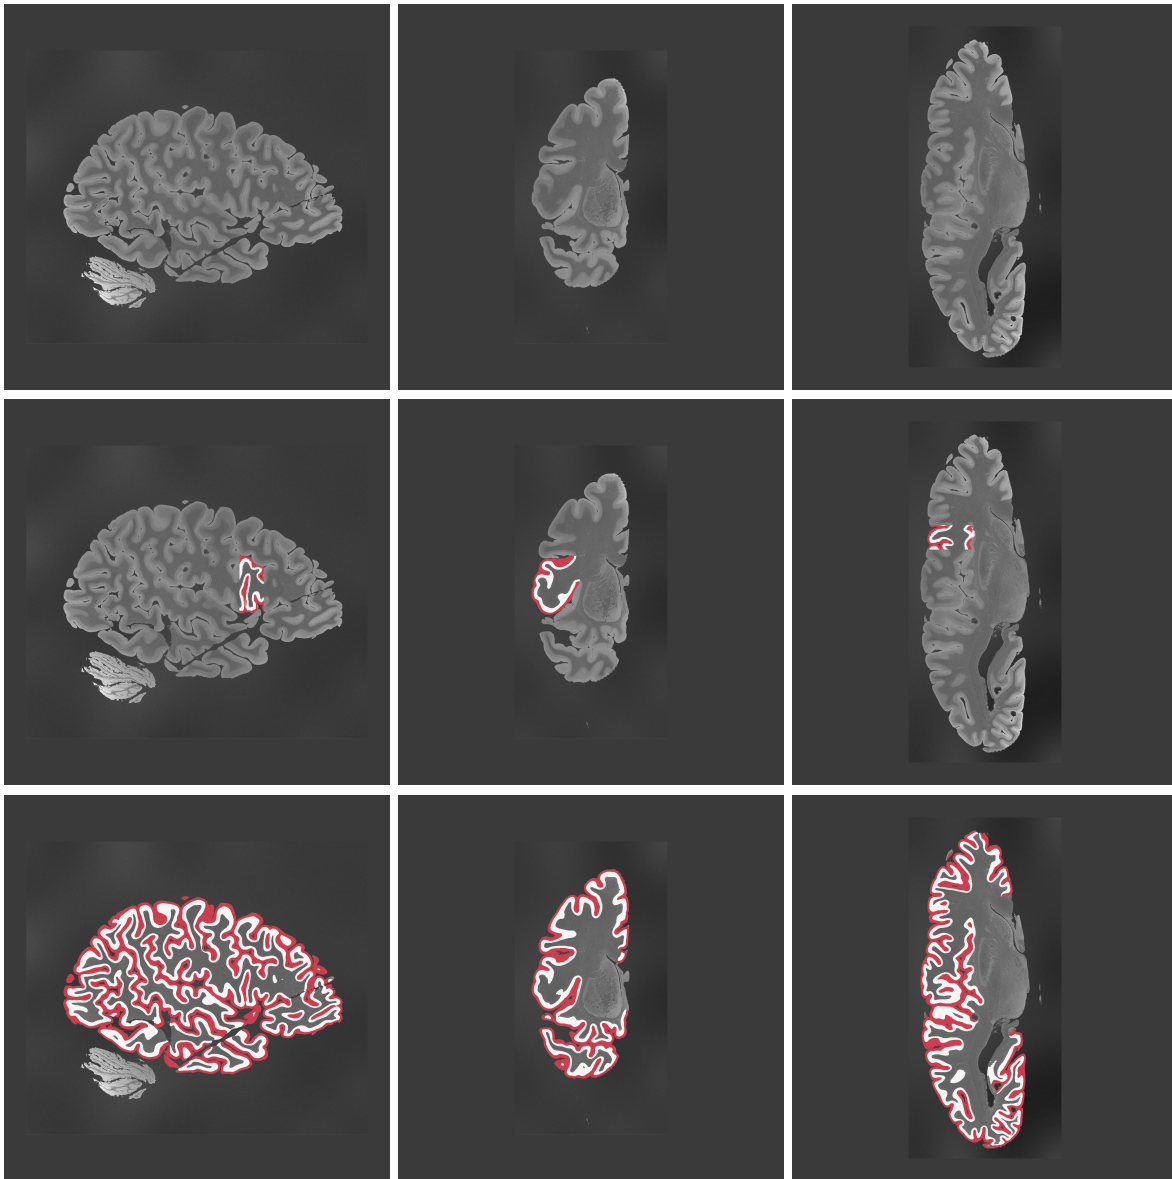Figure 4: 2D slices of case 4 in sagittal, coronal, and axial views.
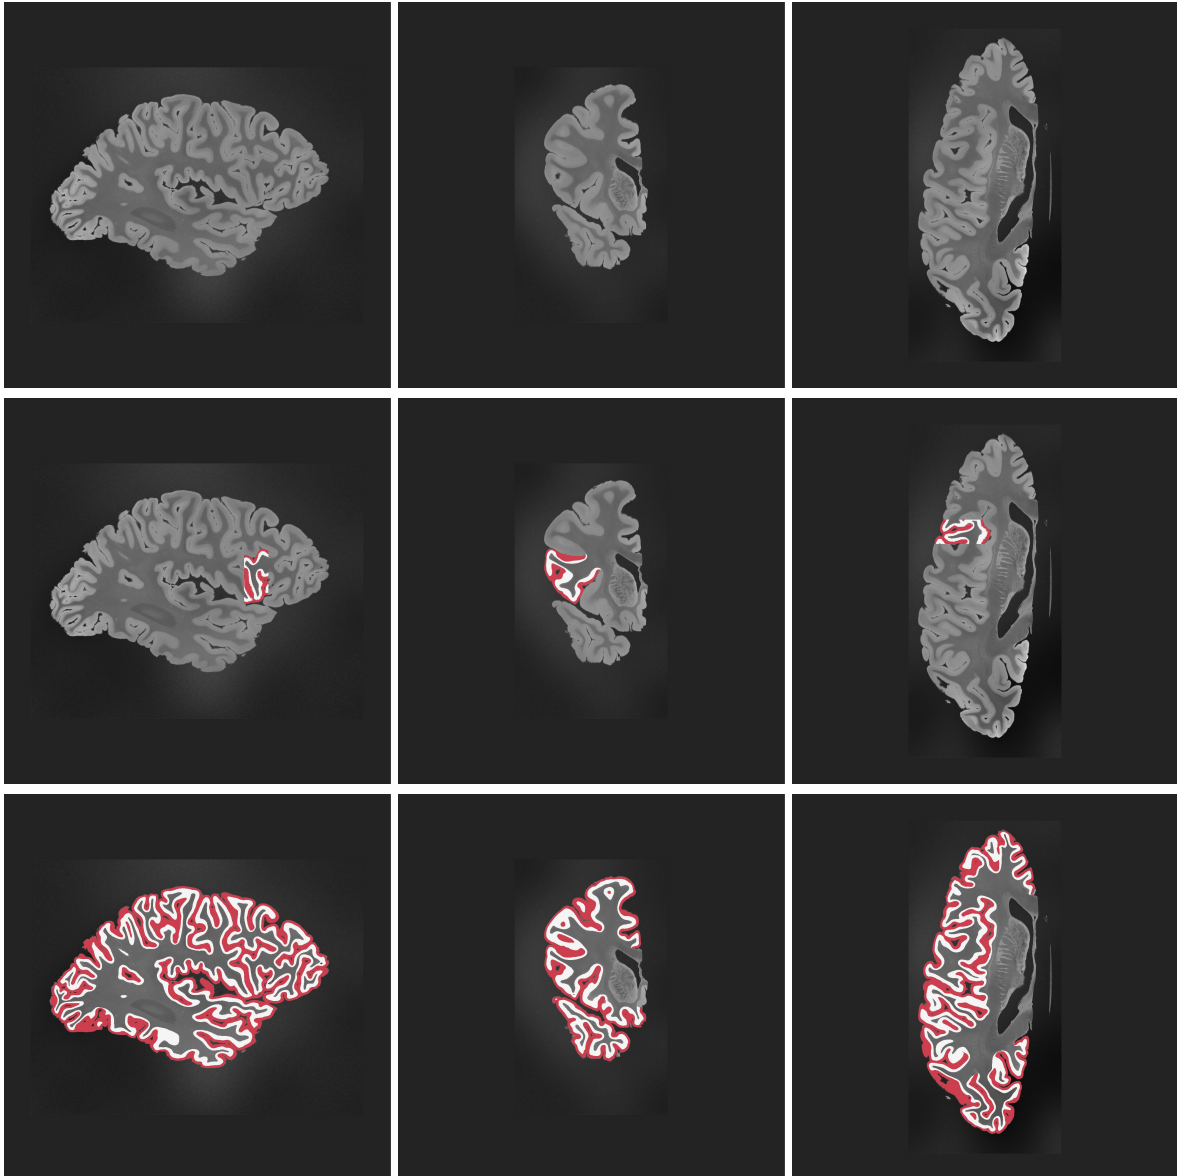
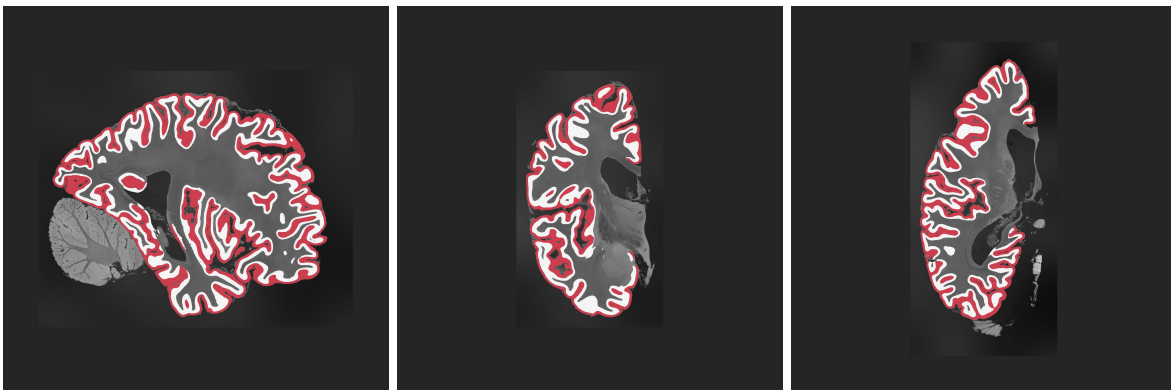Figure 5: 2D slices of case 5 in sagittal, coronal, and axial views.

Figure 6: 2D slices of case 6 in sagittal, coronal, and axial views.

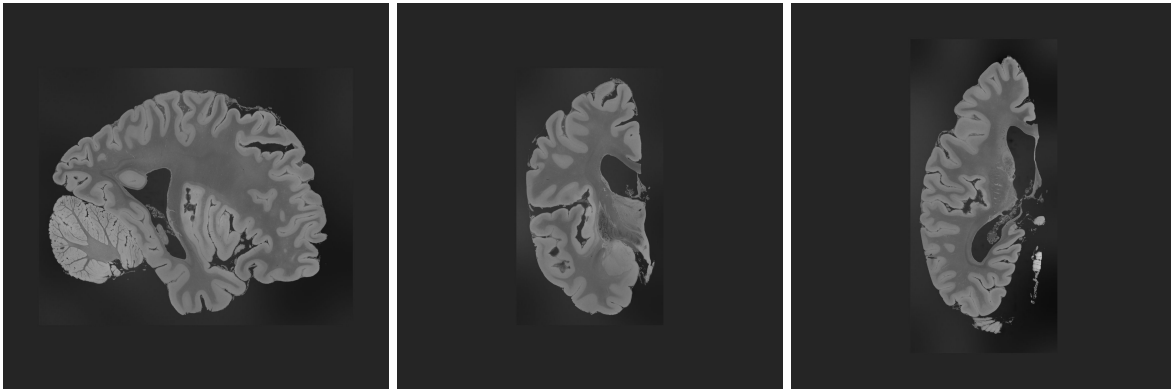Figure 7: 2D slices of case 7 in sagittal, coronal, and axial views.

Figure 8: 2D slices of case 8 in sagittal, coronal, and axial views.

Figure 9: 2D slices of case 9 in sagittal, coronal, and axial views.

Figure 10: 2D slices of case 10 in sagittal, coronal, and axial views.

11

Figure 11: 2D slices of case 11 in sagittal, coronal, and axial views.

Figure 12: 2D slices of case 12 in sagittal, coronal, and axial views.

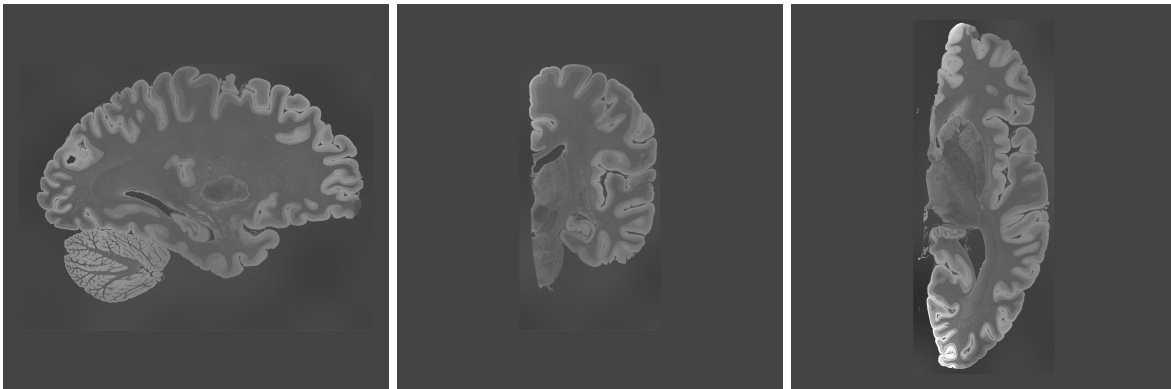Figure 13: 2D slices of case 13 in sagittal, coronal, and axial views.

Figure 14: 2D slices of case 14 in sagittal, coronal, and axial views. This case has not been manually segmented.
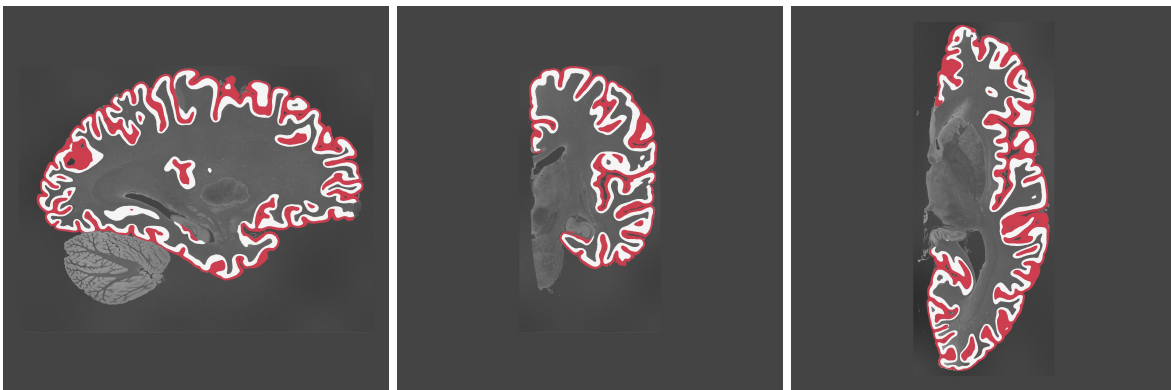
Figure 15: 2D slices of case 15 in sagittal, coronal, and axial views. This case has not been manually segmented.
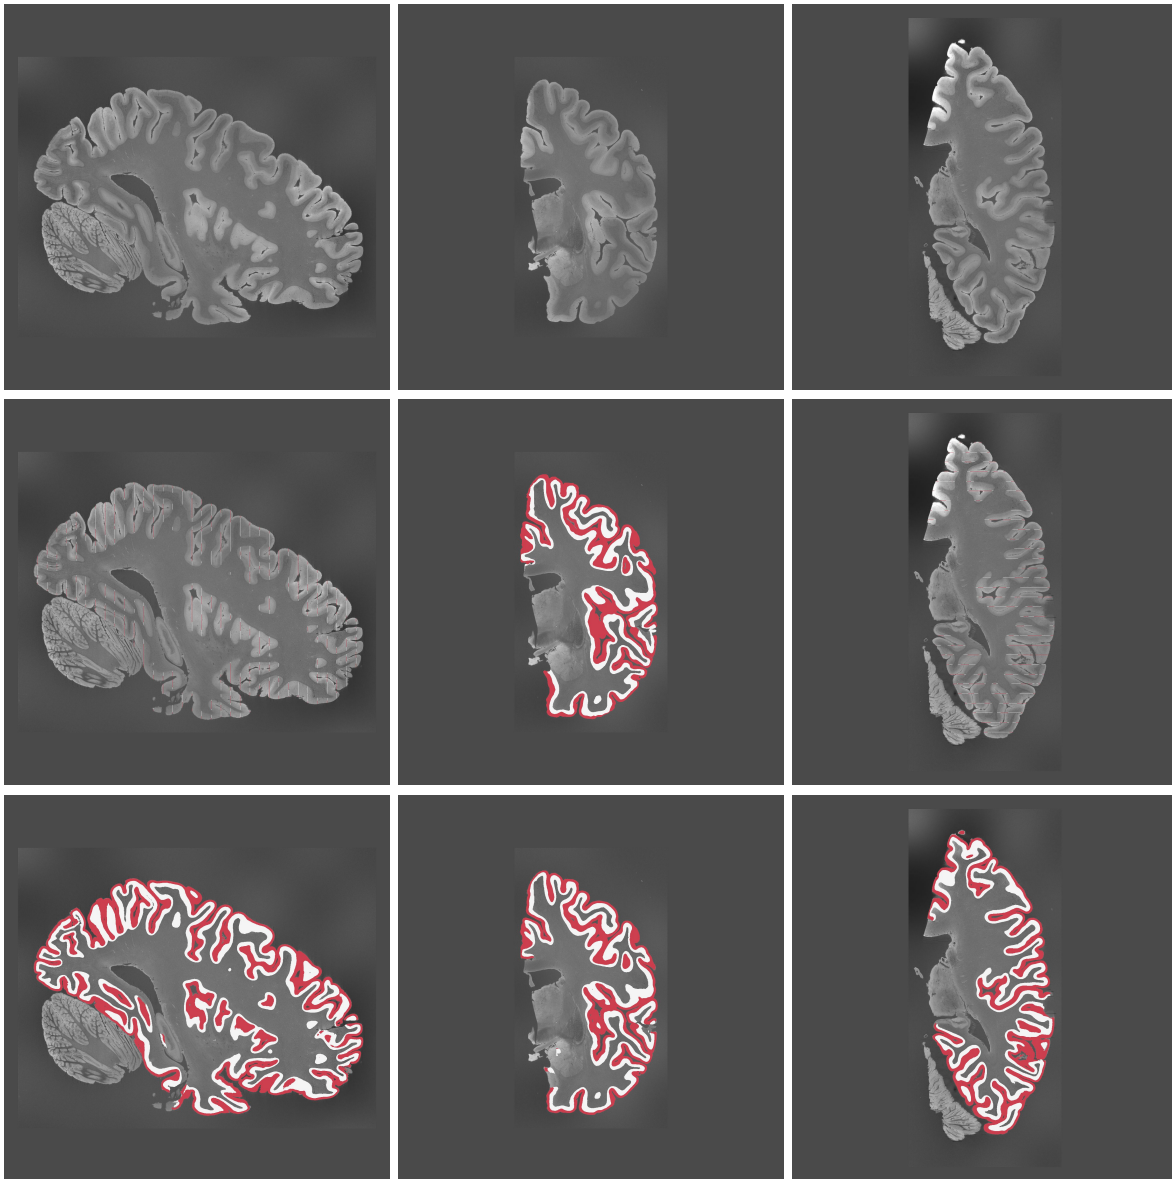
Figure 16: 2D slices of case 16 in sagittal, coronal, and axial views. 29 slices were manually segmented along the coronal view.
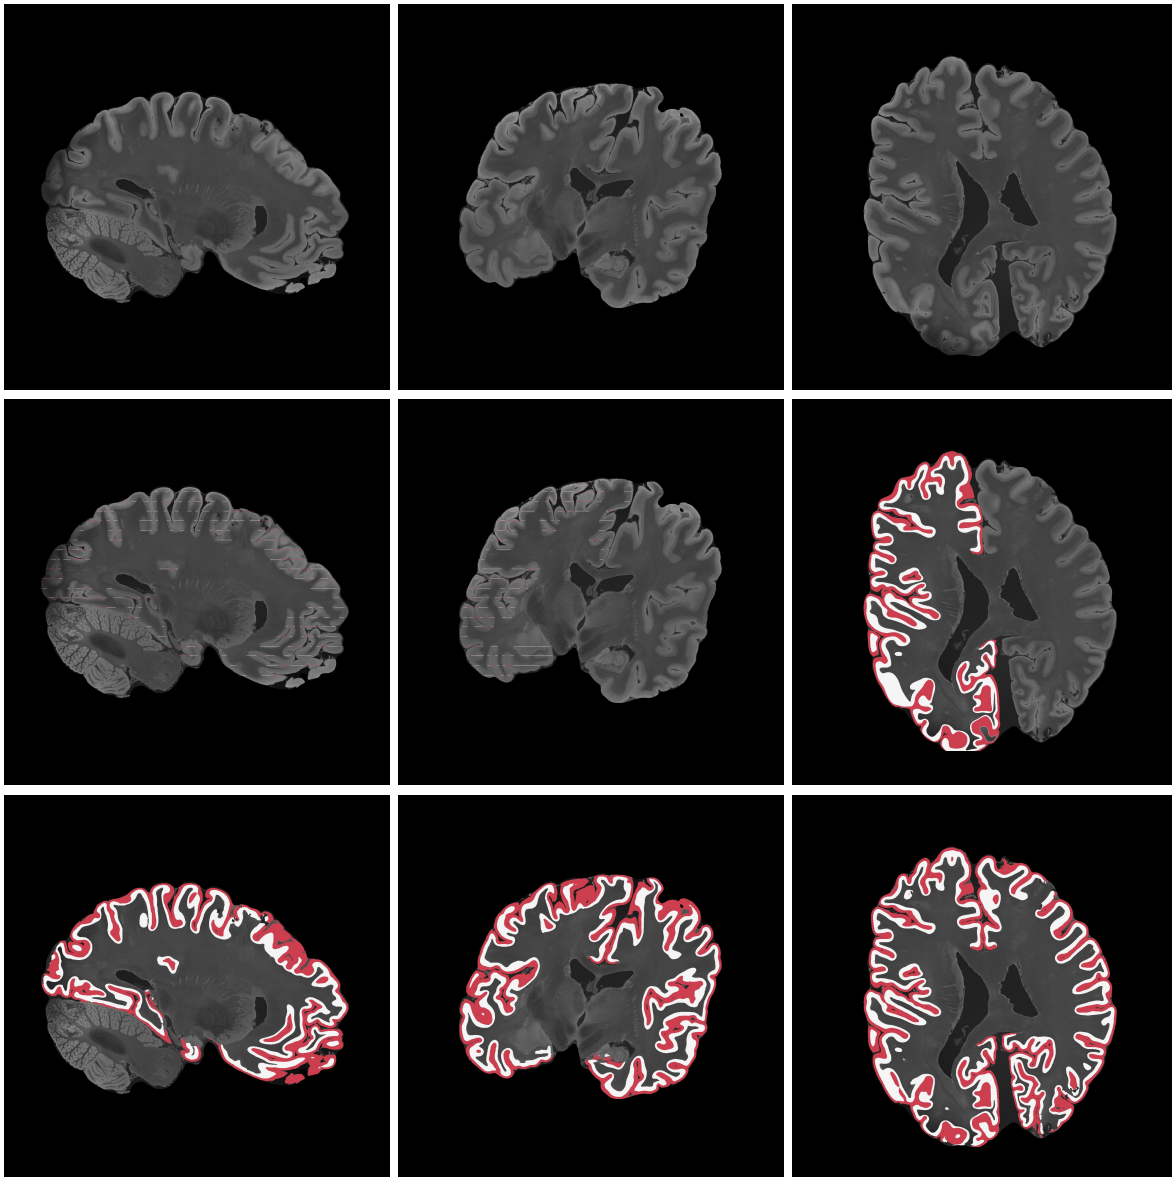
Figure 17: 2D slices of case 17 in sagittal, coronal, and axial views. 21 slices were manually segmented along the axial view.