

Automated olfactory bulb segmentation on high resolutional T2-weighted MRI

Santiago Estrada^{a,b}, Ran Lu^b, Kersten Diers^a, Weiyi Zeng^b, Philipp Ehses^c, Tony Stöcker^{c,d},
Monique M. B. Breteler^{b,e}, Martin Reuter^{a,f,g,*}

^a Image Analysis, German Center for Neurodegenerative Diseases (DZNE), Bonn, Germany

^b Population Health Sciences, German Center for Neurodegenerative Diseases (DZNE), Bonn, Germany

^c MR Physics, German Center for Neurodegenerative Diseases (DZNE), Bonn, Germany

^d Department of Physics and Astronomy, University of Bonn, Germany

^e Institute for Medical Biometry, Informatics and Epidemiology (IMBIE), Faculty of Medicine, University of Bonn, Bonn, Germany

^f A.A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Boston MA, USA

^g Department of Radiology, Harvard Medical School, Boston MA, USA

ARTICLE INFO

Keywords:

Olfactory bulb
Convolutional neural networks
Deep learning
Semantic segmentation

ABSTRACT

The neuroimage analysis community has neglected the automated segmentation of the olfactory bulb (OB) despite its crucial role in olfactory function. The lack of an automatic processing method for the OB can be explained by its challenging properties (small size, location, and poor visibility on traditional MRI scans). Nonetheless, recent advances in MRI acquisition techniques and resolution have allowed raters to generate more reliable manual annotations. Furthermore, the high accuracy of deep learning methods for solving semantic segmentation problems provides us with an option to reliably assess even small structures. In this work, we introduce a novel, fast, and fully automated deep learning pipeline to accurately segment OB tissue on sub-millimeter T2-weighted (T2w) whole-brain MR images. To this end, we designed a three-stage pipeline: (1) Localization of a region containing both OBs using *FastSurferCNN*, (2) Segmentation of OB tissue within the localized region through four independent *AttFastSurferCNN* - a novel deep learning architecture with a self-attention mechanism to improve modeling of contextual information, and (3) Ensemble of the predicted label maps. For this work, both OBs were manually annotated in a total of 620 T2w images for training (n=357) and testing. The OB pipeline exhibits high performance in terms of boundary delineation, OB localization, and volume estimation across a wide range of ages in 203 participants of the Rhineland Study (Dice Score (Dice): 0.852, Volume Similarity (VS): 0.910, and Average Hausdorff Distance (AVD): 0.215 mm). Moreover, it also generalizes to scans of an independent dataset never encountered during training, the Human Connectome Project (HCP), with different acquisition parameters and demographics, evaluated in 30 cases at the native 0.7 mm HCP resolution (Dice: 0.738, VS: 0.790, and AVD: 0.340 mm), and the default 0.8 mm pipeline resolution (Dice: 0.782, VS: 0.858, and AVD: 0.268 mm). We extensively validated our pipeline not only with respect to segmentation accuracy but also to known OB volume effects, where it can sensitively replicate age effects ($\beta = -0.232$, $p < .01$). Furthermore, our method can analyze a 3D volume in less than a minute (GPU) in an end-to-end fashion, providing a validated, efficient, and scalable solution for automatically assessing OB volumes.

1. Introduction

1.1. Motivation

Over the past decades, there has been an increasing awareness to odor function not only as a quality of life indicator (Croy et al. (2014)) but also as a potential biomarker in population studies. Olfactory dysfunction is among the earliest signs of many neurode-

generative disorders, including Alzheimer's and Parkinson's disease (Attems et al. (2014); Doty (2017); Roberts et al. (2016)). Therefore, it is of major interest to gain insights into the anatomical basis of the olfactory pathway in vivo.

New developments in magnetic resonance imaging (MRI) (e.g. field strength, accelerated acquisition schemes, etc.) have allowed the acquisition of high-resolutional (High-Res) MR images, providing an option for reliable assessment of odor-related brain structures, including olfactory bulb (OB). The OB is considered the

* Corresponding author.

E-mail address: martin.reuter@dzne.de (M. Reuter).

<https://doi.org/10.1016/j.neuroimage.2021.118464>.

Received 19 April 2021; Received in revised form 9 July 2021; Accepted 9 August 2021

Available online 10 August 2021.

1053-8119/© 2021 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

most important relay station in the odor pathway, integrating peripheral and central olfactory information. Moreover, OB volume has been associated with olfactory dysfunction in clinical settings (Hummel et al. (2011); Mazal et al. (2016)). However, compared to its central counterparts, i.e. prefrontal cortex, hippocampus, and insular cortex (Dintica et al. (2019); Vassilaki et al. (2017)), OB remains relatively poorly studied, especially in the general population. One reason for that could be the lack of a fully automated segmentation tool for this structure.

Currently, the gold standard for measuring OB volumes is the manual segmentation of T2 weighted (T2w) images—a very expensive and time-consuming process that greatly relies on the raters' expertise. Thus, especially for large population-based studies, automatic segmentation methods are required. However, achieving good accuracy on this small structure is challenging due to its inherent properties: (i) low contrast on T1w scans, (ii) low boundary contrast on T2w images (partial volume effects), (iii) highly sensitivity to noise due to its proximity to the nostrils (e.g. breathing artefacts), (iv) not visible in all subjects (Weiss et al. (2020)), and (v) highly dependent of age (Buschhüter et al. (2008); Hummel et al. (2011, 2015)). So far, those limitations have impeded the wide implementation of any automatic or semi-automatic techniques. Therefore, the introduction of an accurate automated method for segmenting OB is of significant clinical and research interest.

1.2. Olfactory bulb segmentation

Despite the fact, that many studies have analyzed the OB, there is a lack of accurate automatic processing methods for this structure which has been overlooked by many of the standard neuroimage processing frameworks, such as *FreeSurfer* (Fischl et al. (2002)), *BrainSuite* (Shattuck and Leahy (2002)), *SPM* (Friston (2003)), *ANTs* (Avants et al. (2009)), or *FSL* (Jenkinson et al. (2012)). To date, manual delineation is still the predominant approach for accurate quantification of OB volumes. Most groups approximate OB volumes from 1.5T T2w MR scans with a relative low resolution (of 1.5 mm to 2 mm isotropic) (Buschhüter et al. (2008); Hummel et al. (2011, 2015); Seubert et al. (2012)). Recent studies (Joshi et al. (2020); Weiss et al. (2020)) on 3T high-resolution T2w MRI have focused on developing semi-automatic techniques to reduce manual annotations workload but cannot automatically segment the OB. Concurrently to our work, Noothout et al. (2021) proposed an automatic pipeline using fully convolutional neural networks (F-CNNs) to segment the OB on coronal T2w images with an in-plane resolution of 0.47 mm × 0.47 mm and 1 mm thickness. While this method, which is not publicly available at this time, shows promising results in a small dataset (n=21), it is reported to be sensitive to motion artefacts and unseen scenarios (i.e. cases with no apparent OB).

Recently, supervised learning using F-CNNs (Badrinarayanan et al. (2017); Long et al. (2015)) has become the preferred standard in the medical computer vision community for solving semantic segmentation problems when sufficient training data is available (Billot et al. (2020); Dong et al. (2017); Estrada et al. (2020); Henschel et al. (2020); Kamnitsas et al. (2017b); Milletari et al. (2016); Noothout et al. (2021); Ronneberger et al. (2015); Roy et al. (2019b, 2018)). F-CNNs often outperform other traditional methods, as they can learn intrinsic features and integrate global context to resolve local ambiguities in an end-to-end fashion. The most frequently employed network layout for semantic segmentation is the encoder-decoder architecture, i.e. the *UNet* (Ronneberger et al. (2015)). The accuracy of this architecture, however, decreases when segmenting smaller structures (Billot et al. (2020); Estrada et al. (2018); Roy et al. (2018)). This can be due to the more complex shapes (i.e. thinner, irregular boundaries) and visual appearance characteristics in medical images (i.e. less visible and partly occluded). Nonetheless, some of the fault can be attributed to the encoder-decoder layout as it can lead to a

redundant use of information and insufficient encoding of the global contextual information (Fu et al. (2019); Sinha and Dolz (2020)). An accurate understanding of the spatial context is of tremendous importance when segmenting smaller structures as local representation differences between pixels/voxels of a same structure introduce inter-class inconsistencies and affect the recognition accuracy (Fu et al. (2019)). To solve this issue, attention modules have been introduced to improve the understanding of long-range dependencies, not only for semantic segmentation (Fu et al. (2019); Roy et al. (2018); Sinha and Dolz (2020)) but also for other computer vision tasks (Lin et al. (2016, 2017); Vaswani et al. (2017); Zhang et al. (2019)).

In this work, we modify our *FastSurferCNN* (Henschel et al. (2020)) for whole-brain segmentation to focus on the OB. To improve *FastSurferCNN*'s performance for small structures, we suitably included the self-attention mechanism proposed in Zhang et al. (2019) into *FastSurferCNN*; the new deep-learning architecture is termed *AttFastSurferCNN*. *AttFastSurferCNN* promotes attention to spatial information by improving the modeling of local and global-range dependencies. Overall, to segment the OB on high-resolution T2w whole-brain MRI in a fully automatic fashion, we introduce a deep learning pipeline consisting of three stages:

1. **Localization** of a region of interest (ROI) containing the OBs of both hemispheres using a semantic segmentation approach by implementing *FastSurferCNN*; we use the centroid of the predicted region as a center point for cropping a localized volume.
2. **Segmentation** of OB tissue within the localized volume through four *AttFastSurferCNN* with different training condition (four data-splits and data initialization).
3. **Ensemble** stage where the previously generated label maps are averaged and view-aggregated to form a consensual final segmentation.

The presented networks were trained with manual annotations of 357 T2w scans from the Rhineland Study, an ongoing large population-based cohort study (Breteler et al. (2014); Stöcker (2016)). We extensively validated the quality of the individual stages of the pipeline through assessment of segmentation accuracy in an independent unseen heterogeneous in-house dataset ($n = 203$). We showed that our previously introduced *FastSurferCNN* can precisely localize the region containing both OBs and that the proposed *AttFastSurferCNN* can accurately segment the OBs, outperforming other established F-CNNs and accomplishing equivalent results as manual raters. After asserting segmentation accuracy, we validated the soundness of the proposed pipeline in the Rhineland Study with respect to: i) replication of known OB volume effects (e.g. age), ii) stability of volume estimates among variations of the study's T2w sequences, and iii) robustness to scans without an apparent OB. We further assessed generalizability to an unseen externally labeled dataset of 30 subjects from a cohort with different characteristics and acquisition parameters. To the best of our knowledge, our pipeline is the first framework capable of automatically segmenting the OB in a large cohort dataset with high accuracy and reliability. Furthermore, we demonstrated that our method can generalize to different T2w scans with 0.8 mm isotropic resolution. The proposed method is available as an open-source project at: <https://github.com/Deep-MI/olf-bulb-segmentation>.

2. Methodology

2.1. Manual reference standard

Our manual reference standard is based on the annotation of high-resolution (0.8 mm isotropic) T2w MRI from the Rhineland Study. The Rhineland Study (www.rheinland-studie.de) is an ongoing study that enrolls participants aged 30 years and above at baseline from Bonn, Germany. The study is carried out in accordance with the recommendations of the International Council for Harmonisation (ICH) Good Clinical Prac-

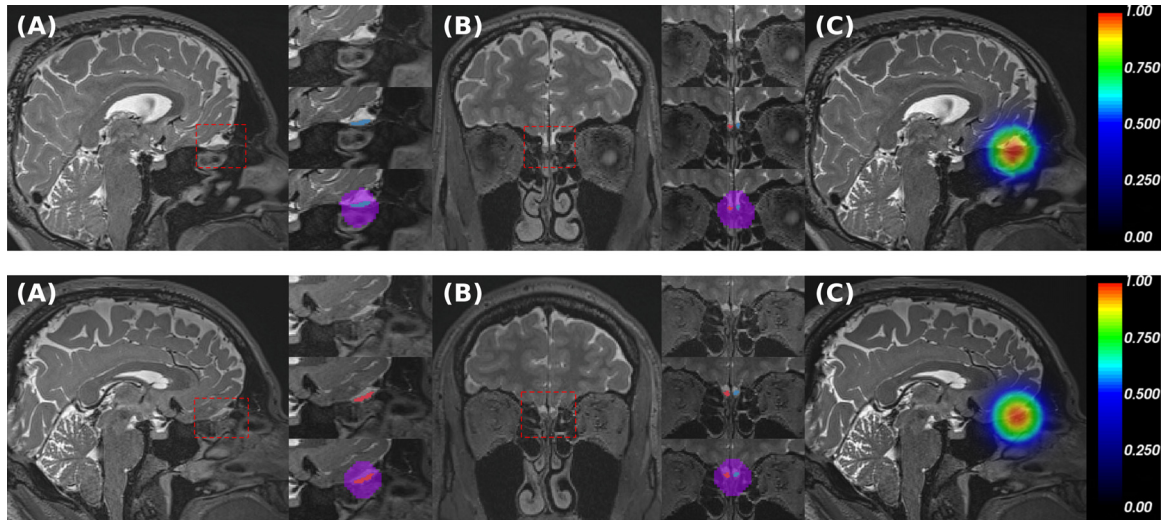


Fig. 1. T2-weighted images and ground truth from two subjects. The red square represents the zoom-in region. A) Sagittal view and labels (blue: left OB, red: right OB, purple: ROI label). B) Coronal view and labels. C) ROI distance map around the centroid of the OB labels. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

tice (GCP) standards (ICH-GCP). Written informed consent was obtained from all participants in accordance with the Declaration of Helsinki.

Manual annotations of the left and right OB were performed by an experienced rater in (unprocessed) T2w images using *Freeview* - a visualization tool of *FreeSurfer* (Fischl (2012); Fischl et al. (2002)). OB is defined as a mostly almond- or spindle-shaped structure symmetrically located at the base of the forebrain (Rombaux et al. (2009)) as seen in Fig. 1, which can be demarcated based on surrounding cerebrospinal fluid and the underlying cribriform plate. The abrupt changes in diameter at the beginning of the olfactory tract in the axial and sagittal views were used as a posterior ending landmark (Wang et al. (2011); Yousem et al. (1998)). In addition, to avoid bias, labeling was blind to participant metadata, e.g. outcomes of the olfactory function and demographics.

For the localization task, we solve a semantic segmentation problem with the goal to segment the forebrain region containing the OBs from both hemispheres (referred to as “region of interest (ROI)”). The ROI label generation is achieved by the following steps: (1) Localization of the mid-point between left and right OB by calculating the centroid (C) of the manual labels. (2) Generation of a distance map by applying a Gaussian distribution around C on a down-sampled 1.6 mm isotropic image, the distance map is defined as: $f(x, y, z) = \mathcal{N}(\mu = C, \delta = 10)$ where x, y and z are voxel coordinates in the down-sampled image. (3) A binary cutoff at $f(x, y, z)/\max(f(x, y, z)) \geq 0.8$ separates ROI and background. The Resulting distance maps and labels are illustrated in Fig. 1.

2.2. Olfactory bulb pipeline

Our proposed deep learning method is aimed at segmenting the OB on high-resolution T2w whole-brain MRI. This task presents the challenge of a high-class imbalance between foreground and background ($\approx 1 : 10^6$). A reduction in the spatial size of the input can partially alleviate the problem by cropping the background and by focusing the background information on relevant regions in close proximity to the OBs. This, furthermore, reduces computational and memory requirements during training and inference. Following this direction, we designed a fully automated pipeline for OB tissue segmentation as depicted in Fig. 2.

The proposed pipeline consists of three stages: (1) In order to remove most of the unnecessary background we first train *FastSurferCNN* (Henschel et al. (2020)) with a down-sampled 1.6 mm isotropic image to provide a quick segmentation of the forebrain region contain-

ing both OBs (localization network). This segmentation is only used to compute a centroid coordinate of the region of interest. A final localized volume (at 0.8 mm isotropic, $96 \times 96 \times 96$ voxels), centered at this coordinate, is cropped or resampled from the input image. By default the pipeline resamples deviating resolutions to 0.8 mm isotropic, unless the user specifies to use the native resolution instead. (2) Afterwards, the OB tissue is segmented within this cropped volume by four *AttFastSurferCNN*s with different training conditions (four data-splits and data initialization). (3) Finally, the ensemble segmentation is composed by averaging the predicted label maps; the implemented ensemble approach ensures that only voxels with high agreement among models are selected and also reduces variance due to network initialization. Furthermore, since right and left OB were combined as one structure during segmentation, they are split retrospectively in an independent post-processing step.

2.2.1. Region of interest (ROI) localization network - *FastSurferCNN*

To localize the ROI as a semantic segmentation task, we employ *FastSurferCNN* (Henschel et al. (2020)) as it outperformed other commonly used encoder-decoder architectures, i.e. *SDNet* (Roy et al. (2017)) and *QuickNat* (Roy et al. (2019b)), on the difficult task of whole-brain segmentation. *FastSurferCNN* consists of three 2D F-CNNs operating on different anatomical views (coronal, axial, and sagittal) and a final view-aggregation stage. In brief, all F-CNNs follow the same layout of four competitive-dense blocks (CDB) for the encoder and decoder path separated by a bottleneck block. The use of CDB reduces the number of learnable parameters by replacing the typical concatenation units inside dense-connections with maxout activations (Goodfellow et al. (2013); Huang et al. (2017)). The maxout activation induces competition between feature maps by computing the maximum at each spatial location, thus improving the feature selectivity (Liao and Carneiro (2017)) and boosting the learning of fine-grained structures (Estrada et al. (2018, 2020)). Furthermore, *FastSurferCNN* utilizes a multi-slice input approach by stacking preceding slices, current, and succeeding slices for segmenting only the middle slice, which in turn increases the spatial information aggregation in a 2D network by improving the local neighborhood awareness (Henschel et al. (2020)).

In this work, we slightly modified *FastSurferCNN* by adjusting the view-aggregation step to a normal unweighted average. Since the ROI label is not lateralized, there is no need to increase attention to any particular anatomical view. Furthermore, the prior downsampling of the input scan (to isotropic 1.6 mm) allows a reduction of the multi-slice in-

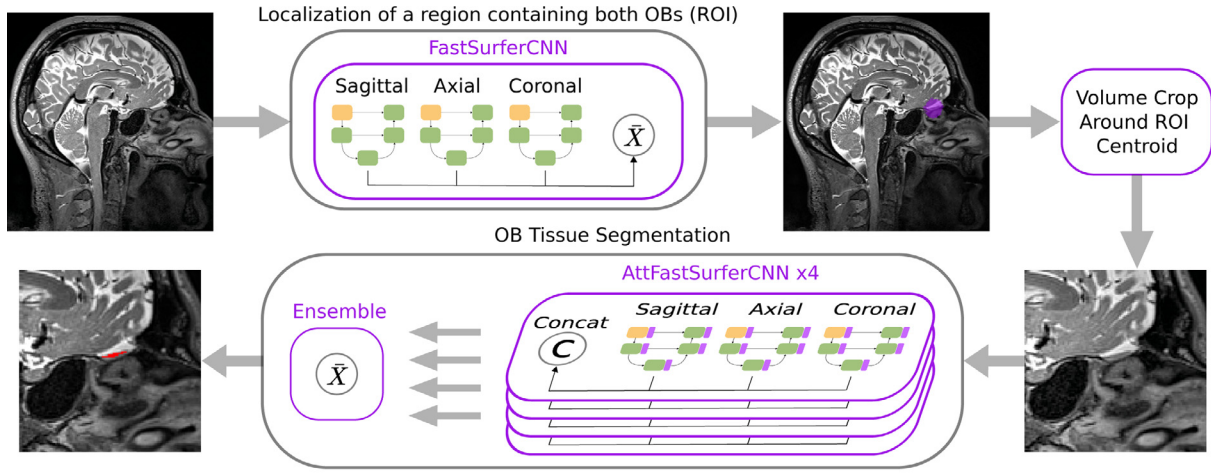


Fig. 2. Proposed pipeline for OB segmentation. The pipeline is divided into three stages: First, localization of a region of interest containing the left and right OB. Then, OB tissue segmentation within the localized volume, and finally, an ensemble of predicted label maps.

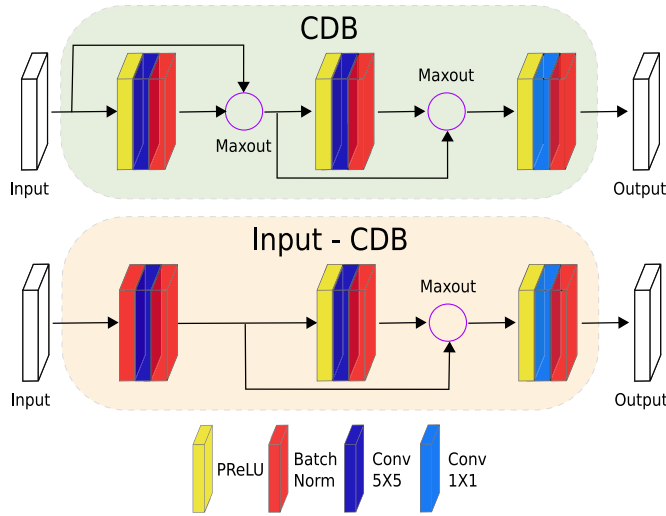


Fig. 3. Competitive Dense Blocks (CDB) configuration. Each block is composed of three sequences of parametric rectified linear unit (PReLU), convolution (Conv) and batch normalization (BN) (bottom) with exception of the very first encoder block (top). In the first block, the PReLU is replaced with a BN to normalize the raw inputs.

put image from 7 to 3 consecutive slices while retaining approximately the same field of view. In terms of the CDB blocks, the three configuration sequences of a parametric rectified linear unit (PReLU), convolution (Conv)(64 filters), and batch normalization (BN) are maintained (Fig. 3 top) as well as the exception for the very first encoder block. In the first block, the first PReLU is replaced with a BN to normalize the raw inputs (Fig. 3 bottom).

2.2.2. OB Segmentation network - AttFastSurferCNN

To accurately segment the OB, we introduce *AttFastSurferCNN* a new deep learning architecture that boosts the attention to spatial information. We implemented *AttFastSurferCNN* by suitably including the self-attention mechanism proposed by Zhang et al. (2019) into *FastSurferCNN* (Henschel et al. (2020)). The self-attention module was included after each competitive-dense block (CDB), as shown in Fig. 4, thus improving the modeling of contextual information. Furthermore, in order to take full advantage of the multi-scale attention maps (Fu et al. (2019); Sinha and Dolz (2020)) and to prevent information loss from the unpooling layers (Estrada et al. (2018)), we replaced the maxout activation

units between the finer feature maps from long-range skip connections and the coarser feature maps from the unpooling path with an element-wise sum.

The implemented self-attention layer is illustrated in Fig. 5. Let us denote the CDB output feature map as $F_{CDB} \in \mathbb{R}^{C \times H \times W}$, where C, H, W represent the channel, height, and width dimensions respectively. First, the F_{CDB} is fit into two 1×1 convolutional layers to reshape the channels to a size of $C/8$ and create two new features maps (F_a and F_b). Reducing the number of channels drastically diminishes memory requirements without a significant performance loss (Zhang et al. (2019)). Subsequently, the feature maps are flattened to a shape of $(C/8) \times (N)$, where $N = H \times W$ are the number of pixels. Afterwards, an attention map (S) is created by applying a softmax layer into the output of a matrix multiplication between F_a^T and F_b . Thus $S \in \mathbb{R}^{N \times N}$ is defined as:

$$s_{j,i} = \exp(F_a^T \cdot F_b) / \sum_{i=1}^N \exp(F_a^T \cdot F_{b_i}) \quad (1)$$

where $s_{j,i}$ indicates the extend to which the i_{th} position impacts the j_{th} position. Before applying S , the F_{CDB} features are fed into a 1×1 convolutional layer and a new feature map $F_c \in \mathbb{R}^{C \times H \times W}$ is generated and reshaped to $\mathbb{R}^{C \times N}$. Afterwards, a matrix multiplication is performed between the transpose of S and F_c and the results reshaped to the original size $\mathbb{R}^{C \times H \times W}$. Finally, the self-attention output (F_{att}) is formulated as follows:

$$F_{att} = \alpha(S^T \cdot F_c) + F_{CDB} \quad (2)$$

where α is a learnable scalar parameter initialized with 0. The introduction of α allows the network to first focus on the local information which is an easier task and gradually increases the importance of non-local dependencies which is a harder task (Zhang et al. (2019)). We additionally normalize F_{att} thus guaranteeing a normalized input to the other CDB blocks. A normalized input improves convergence (Liao and Carneiro (2016)) and increases the exploratory span of the created sub-networks when using a maxout activation (Liao and Carneiro (2017)). In summary, the implemented spatial attention module improves the modelling of local and global-range dependencies, which in turn increases semantic consistency.

In brief, *AttFastSurferCNN* is a multi-network approach of three 2D F-CNNs operating on different anatomical views (coronal, sagittal and, axial). All three F-CNNs contain the self-attention layers following the aforementioned layout (Fig. 4). Within *AttFastSurferCNN* the CDB blocks maintain the configuration from Section 2.2.1 except for the 5×5 convolutions that are modified to a smaller kernel size of 3×3 . Furthermore, the multi-slice input approach from *FastSurferCNN* (Henschel et al. (2020)) is maintained and a stack of three consec-

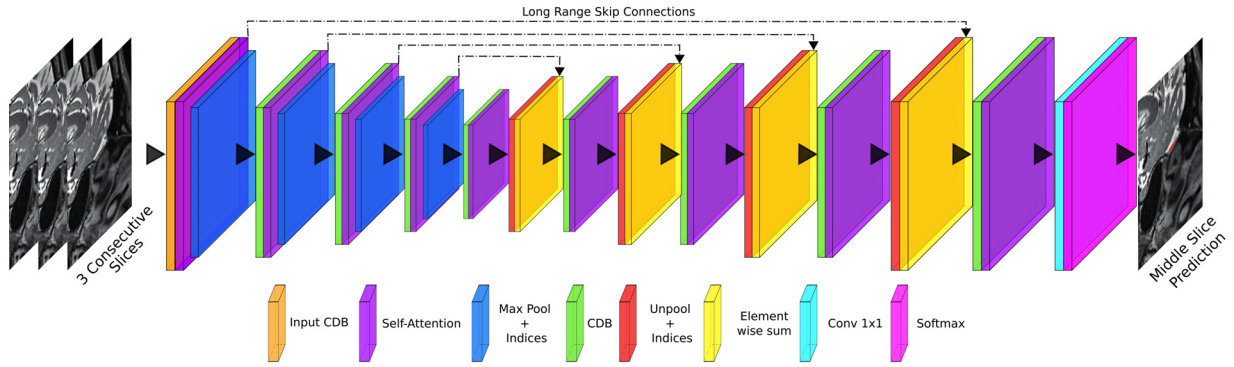


Fig. 4. *AttFastSurferCNN* network architecture. The network consists of four competitive dense blocks (CDB) in the encoder and decoder part, separated by a bottleneck layer. After each CDB a self-attention module is added. The CDB configuration is illustrated in Fig. 3.

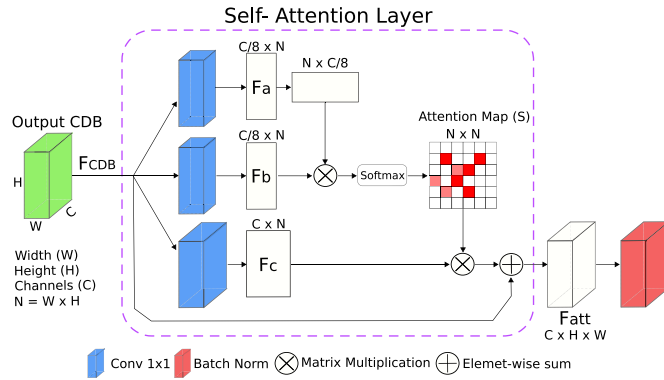


Fig. 5. Implemented Self-Attention module within *AttFastSurferCNN*.

utive slices are passed as input. In the following section, the ensemble of different segmentation predictions will be explained in detail.

2.2.3. OB Segmentation ensemble

One widely used method to assess the optimal model of CNNs trained with different data-splits is cross-validation. Cross-validation jointly evaluates performance on different data-splits and the model with the maximal test-set performance is selected as the winner. This approach, however, can limit generalizability as the data-splits used for training the best performer can be biased towards the selected test-set. Recently, the combination of different CNN model outputs has been shown to improve the prediction performance and reduce the CNN's intrinsic variance (Ju et al. (2018)). As a consequence, we propose to ensemble the prediction of four *AttFastSurferCNN*s trained with different data-splits, ensuring that only OB voxels with a high inter-model agreement are segmented, and thus reducing the bias to any particular data division. To ensure that all networks have a comparable OB knowledge: i) training was done under the same learning conditions (i.e. number of epochs, batch size, loss function, a learning rate scheduler, etc.), ii) training data was divided into four data-splits balanced for age and sex, and iii) the data-splits were treated in a leave-one-out fashion. Finally, the ensemble is constructed by an unweighted average as the output of models with comparable performance is merged (He et al. (2016); Ju et al. (2018); Kamnitsas et al. (2017a); Szegedy et al. (2015)). Intuitively, the proposed ensemble approach can be seen as four different raters with similar experience taught by the same instructor and the consensus among the raters gives the final decision. It is important to note, that in our specific approach the final ensemble prediction is created by averaging twelve different models as each *AttFastSurferCNN* contains three 2D F-CNNs for the three different anatomical views (axial, coronal and, sagittal). Therefore, our ensemble approach also includes the advan-

tages of view-aggregation where a voxel prediction is regularized by considering spatial information from multi-views (Estrada et al. (2020); Henschel et al. (2020); Roy et al. (2019b)). We furthermore analyzed the impact of the ensemble approach by comparing directly with standalone data-splits.

2.2.4. Model learning

All F-CNN models for localization and segmentation were implemented in PyTorch (Paszke et al. (2017)) using a docker container (Merkel (2014)). Independent models for axial, coronal, and sagittal views were trained for 40 epochs with a batch size of 16 using two NVIDIA Tesla V100 GPU with 32 GB RAM, and a Adam optimizer (Kingma and Ba (2015)) with a step decay scheduler that decreases the learning rate (lr) by 95% every 5 epochs (initial lr = 0.01, constant weight decay = 10^{-04} (Loshchilov and Hutter (2019)), betas=(0.9, 0.999), eps= 10^{-08}). The networks were trained by optimizing a composed loss function of focal loss (Lin et al. (2017)) and dice loss (Milletari et al. (2016)). The focal loss addresses the class imbalance by modifying the standard cross-entropy loss such that lower importance is given to the well-classified pixels. On the other hand, the dice loss is a more robust loss to handle data imbalance (Sudre et al. (2017b)) as it is based on the Dice score, an overlay similarity index that reflects both size and localization agreement. Therefore, our proposed composed loss function is formulated as:

$$Loss = \underbrace{- \sum_x w(x)(1 - p_l(x))^\gamma g_l(x) \log(p_l(x))}_{\text{Weighted Focal Loss}} - \underbrace{\frac{2 \sum_x p_l(x)g_l(x)}{\sum_x p_l^2(x) + \sum_x g_l^2(x)}}_{\text{Dice Loss}} \quad (3)$$

where $p_l(x)$ is the predicted probability at pixel x to belong to a class l , and $g_l(x)$ is the pixel ground truth class. For the weighted focal loss, γ was set to 2 and the pixel weight scheme ($w(x)$) proposed by Roy et al. (2019b) was used to improve segmentation performance along anatomical boundaries. We additionally included online data augmentation to address two challenges: 1) spatial variations due to head position and image cropping, and 2) intensity inhomogeneities due to scan parameters and movement artefacts (e.g. eyes and breathing). The first problem was tackled by applying random spatial transformations (translation, rotation, and global scaling) on the input images. It is important to notice that spatial augmentations were done in a full image for the segmentation models before cropping, therefore eliminating the intrinsic padding noise when interpolating cropped images. For the second issue, we improved the network's robustness to intensity variations by performing random bias field (Sudre et al. (2017a)) and blur transformations. To maintain consistency between neighboring slices, intensity transformations were performed on a subject level (whole volume) using TorchIO (Pérez-García et al. (2020)).

Table 1

Summary of the datasets, number of subjects, T2 protocol and usage for each of the validation experiments.

Usage	Dataset Name	T2 protocol	Subjects	Cohort
Manual annotation	In-house train-set	$T2w^a$	31	Rhineland Study
reproducibility (E1)	In-house test-set		19	
Pipeline performance (E2)	In-house train-set	$T2w^a$	357	
	In-house test-set		203	
Age and sex effect	In-house test-set	$T2w^a$	203	
sensitivity (E3)				
No apparent OB (E4)	No-OB set	$T2w^a$	12	
Sequence Stability (E5)	Stability set	$T2w^a, T2w^b$	109	
Generalizability (E6)	HCP dataset	$T2w^{hcp}$	30	Human Connectome Project(HCP)

2.3. MRI Data

MRI scans from the Rhineland Study were collected at two different sites both with identical 3T Siemens MAGNETOM Prisma MRI scanners (Siemens Healthcare, Erlangen, Germany) equipped with 64-channel head-neck coils. The 0.8 mm isotropic T2-weighted 3D Turbo-Spin-Echo (TSE) sequence uses variable flip angles (Busse et al. (2008)) as well as elliptical sampling (Mugler III (2014)) and parallel imaging (PI) (Griswold et al. (2002)) for faster imaging. For this work, two T2w sequences from the Rhineland Study were considered (referred to as $T2w^a$ (original protocol) and $T2w^b$). Common sequence parameters are as follows: repetition time (TR) = 2800 ms, echo time (TE) = 4405 ms, phase-encoding direction: Anterior > Posterior, matrix size = $320 \times 320 \times 224$. The following parameters differ between protocols: PI acceleration factor: a) 3x1; b) 2x1, PI reference scan: a) integrated; b) external, acquisition time: a) 3:57 min; b) 4:47 min. Note, care was taken to preserve the image contrast between versions.

For the training and testing of our pipeline, data from the first 572 participants from the Rhineland Study with a T2w scan was used (referred to as "in-house dataset"). All 572 MRI scans were manually annotated following Section 2.1. During the creation of the in-house dataset, a group of 12 subjects was separated into another subset (referred to as "no-OB dataset") as these cases were flagged with no visible OB. Subjects without an apparent OB had been reported previously (Weiss et al. (2020)). Consequently, the no-OB cases were used to evaluate the automated method's robustness to an unseen extreme scenario. The remaining sample (n=560) presents a mean age of 53.83 years (range 30 to 87), a mean OB volume of 54.05 mm^3 (range 12.80 to 111.10 mm^3), and 56.8% of subjects are women. We initially divided the in-house dataset into a training (n=357) and testing (n=203) set. For each subset subjects were randomly selected from sex and age strata to ensure a balanced population distribution. Training data was further split into four groups with the same stratification scheme as before. For a detailed description of the population characteristics of all the aforementioned subsets see Appendix Table 1.

Additionally, another subset of the Rhineland Study was selected to evaluate the prediction stability across T2w sequences, as the proposed pipeline was trained only with $T2w^a$ scans. As part of the quality assurance workflow in the Rhineland Study before updating a sequence, new incoming subjects are scanned in the same session with both versions for a period of time. After the acquisition reliability is assured the study protocol is updated. Therefore, we selected a group of subjects containing both $T2w^a$ and $T2w^b$ scans (referred to as "stability dataset", n=109).

Finally, we used the publicly available Human Connectome Project (HCP) dataset (Van Essen et al. (2012)) to test the generalizability of our method as it contains high-resolution T2w MR images. A subset of 30 random subjects equally distributed between age categories (22–25, 26–30, and 31–35) was selected. The HCP scans were resampled from isotropic 0.7 mm native resolution to 0.8 mm network input resolution. Finally, manual labels were created for both resolu-

tions using the protocol previously described. HCP data is available at: <https://www.humanconnectome.org/study/hcp-young-adult>.

2.4. Evaluation metrics

For assessing the segmentation similarity between the predicted label maps and the ground truth, we computed metrics aimed at evaluating different properties: spatial overlap, spatial distance, and volume similarity. We first assessed the spatial overlap as it provides both size and localization consensus by computing the Dice similarity coefficient (Dice), which is a common metric used for validating semantic segmentation performance. Let G (ground truth) and P (prediction) denote binary label maps; the Dice similarity coefficient is mathematically expressed as

$$Dice = \frac{2 \cdot |G \cap P|}{|G| + |P|} \quad (4)$$

where $|G|$ and $|P|$ represent the number of elements in each label map, and $|G \cap P|$ the number of common elements, therefore, the Dice ranges from 0 to 1 and a higher Dice represents a better agreement. However, Dice scores can be drastically affected by small spatial shifts when evaluating small and elongated structures such as the OB (Billot et al. (2020); Taha and Hanbury (2015)). Spatial distance-based metrics such as Hausdorff Distance (HD) are widely used for assessing performance in small structures as they evaluate the quality of segmentation boundaries. In this work, we used the Average Hausdorff Distance (AVD), an HD variation less sensitive to outliers. AVD is defined as

$$AVD(G, P) = \max(\frac{1}{|G|} \sum_{g \in G} \min_{p \in P} d(g, p), \frac{1}{|P|} \sum_{p \in P} \min_{g \in G} d(p, g)) \quad (5)$$

where d is the Euclidean distance. In contrast to the Dice, AVD is a dissimilarity measurement so a smaller AVD indicates a better boundary delineation with a value of zero being the minimum (perfect alignment). Furthermore, as the OB volumes are usually the desired marker for downstream analysis, we computed a volume-based metric, the volume similarity (VS (Taha and Hanbury (2015))), defined as

$$VS = 1 - \frac{||G| - |P||}{|G| + |P|} \quad (6)$$

While VS is similar to Dice, it does not take into account segmentations overlap and can have its maximum value even when the overlap is zero. In consequence, VS is not used for the localization marker and replaced with localization distance (R), a metric more suitable to assess the accuracy of the centroid coordinate created in this stage. Let p and g be the centroid coordinates of the predicted and ground truth label maps, respectively. The localization distance (R) is calculated as follows

$$R(p, q) = \sqrt{(p_x - g_x)^2 + (p_y - g_y)^2 + (p_z - g_z)^2} \quad (7)$$

Similar to AVD, a smaller distance indicates improved localization accuracy. Finally, to benchmark performance of various F-CNN models we first ranked the models performance for each metric individually and then computed an overall rank as the geometric mean of the model's rankings.

3. Experiments and results

In this section, we present six experiments with the aim to thoroughly validate our OB tissue segmentation pipeline. To properly assess the pipeline's performance as a whole, input images to the segmentation stage were pre-processed by the localization stage. Additionally, to ensure that all experiments were carried out under the same testing conditions: All inference analyses were evaluated in a docker container with a 12 GB NVIDIA Titan V GPU (a widely available consumer card). It is important to note, that the pipeline can also run on the CPU.

(E1) We evaluated the OB manual annotations reliability by an inter and intra-rater reproducibility analysis. (E2) We evaluated the performance of each stage of the pipeline against an unseen test-set. We additionally benchmarked the proposed *AttFastSurferCNN* with state-of-the-art F-CNNs and compared the accuracy of one *AttFastSurferCNN* against the proposed ensemble approach of merging four *AttFastSurferCNN* with different training-data conditions. (E3) We assessed the sensitivity of the proposed pipeline to replicate known OB volume effects with respect to age and sex on the test-set against manual labels and benchmark networks. (E4) We evaluated the robustness of the automated method to an extreme and real scenario of cases without an apparent OB. (E5) We tested the stability of the proposed pipeline to variations in acquisition parameters of a T2w sequence. Finally (E6), we assessed the generalizability of our method to different population demographics on the publicly available HCP dataset (Van Essen et al. (2012)). A summary of the data needed for each of the experiments is presented in Table 1.

3.1. Manual annotation reproducibility (E1)

To the best of our knowledge, there is no automatic method for detecting and delineating the OB. Therefore, manually annotations are considered the gold standard. As our approach is based on supervised learning, its performance is limited by the quality of the manual annotations. As a consequence, to assess the consistency of the labels created by our main rater, we conducted intra-rater and inter-rater variability experiments.

Fifty random subjects from the in-house dataset were selected. Afterwards cases were manually annotated twice (see Section 2.1), once by our main rater who had already segmented the cases and once by a second rater trained by our main rater. To remove bias and avoid overestimating performance, raters were blind to the scans' identification; furthermore, the main rater's second segmentations were done with a time gap of two months, and finally, the scans used for training the second rater were not included in the experiment. We assessed intra-rater variability by computing the similarity between the two sets of segmentations of the main rater. Inter-rater variability was estimated by comparing the segmentation agreement between the main rater's first annotations and second rater's annotations.

In Fig. 6, we present the similarity scores for total OB (left and right combined) in the fifty subjects used for this experiment as well as significance level indicators (paired two-sided Wilcoxon signed-rank test (Wilcoxon (1992))). We observed that our main rater has a high agreement between labeling sessions (Average : $Dice = 0.9399$, $VS = 0.9811$, $AVD = 0.0976$ mm). Inter-rater scores (Average : $Dice = 0.8211$, $VS = 0.9497$, $AVD = 0.2446$ mm) are significantly lower, however, still yield comparable results with other small brain structures inter-rater-scores (Billot et al. (2020)). These similarity scores put the results of the next section into context where the inter-rater-scores can be seen as the lower-bound of performance and intra-rater-scores as the ideal performance of the automated method.

3.2. Pipeline performance (E2)

In this section, we benchmarked and evaluated the accuracy of each stage of the pipeline in a completely separate unseen test-set. All implemented networks were trained using the scheme mentioned in

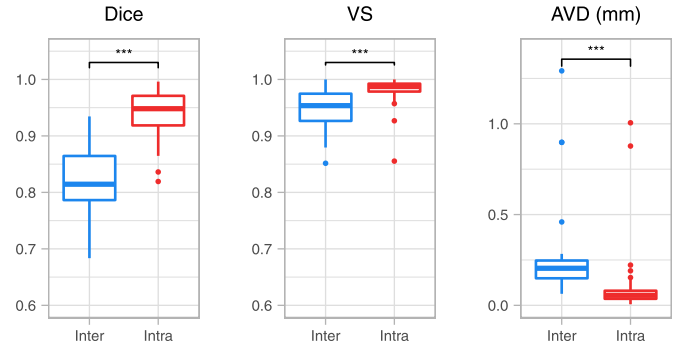


Fig. 6. Segmentation similarity scores for total OB comparing intra-rater vs. inter-rater variability, as well as significance level indicators (paired two-sided Wilcoxon signed-rank). Significance: *** $p < 0.001$.

Section 2.2.4 and data-splits introduced in Section 2.3 were treated in a leave-one-out fashion (e.g. model 1: splits 2, 3, and 4 were used for training, and split 1 was used for validation).

3.2.1. ROI Localization

For evaluating the ability of *FastSurferCNN* to localize the OB ROI in a down-sampled whole-brain image, we trained *FastSurferCNN* from scratch using the four data-splits from the in-house train-set in a leave-one-out cross-validation approach. To ensure good performance and reduce initialization variance, each data-split was trained four times, and the best weights per split were chosen based on the performance in the validation-set. Finally, the model with the highest overall rank of the three evaluation metrics (Dice, AVD, R) in the test-set was selected and incorporated into the pipeline's localization stage.

We observed that all *FastSurferCNN* models have comparable results when segmenting the ROI (Average : $Dice \approx 0.83$, $AVD \approx 0.4$ mm) with model 4 outperforming models 2 and 3 with statistical significance as illustrated in Appendix Fig. A1. However, the small shifts on the predicted label maps did not affect the coordinates from the computed centroid as all models have similar performance ($R \approx 2.08$ mm); hence, any of the trained *FastSurferCNN*s could be used for localizing the ROI for cropping. However, we selected the *FastSurferCNN* trained with data from splits 1, 2, and 3 (model 4) as it has the highest overall rank and outperforms the other versions.

3.2.2. OB Tissue segmentation

To show a proof-of-concept for our proposed *AttFastSurferCNN* in the more difficult task of OB tissue segmentation, we benchmarked our network against state-of-the-art segmentation 2D F-CNNs used for neuro-imaging such as *FastSurferCNN* (Henschel et al. (2020)), *UNet* (Ronneberger et al. (2015)), and *QuickNat* (Roy et al. (2019b)). Additionally, we compared our *AttFastSurferCNN* against 3D networks such as *3D-UNet* (Çiçek et al. (2016)) and *3D-FastSurferCNN*, a naive 3D implementation of *FastSurferCNN* by replacing 2D operations for 3D ones. To permit a fair comparison, all benchmark networks followed the same architecture of four encoder blocks, four decoders blocks, and one bottleneck block as illustrated in Fig. 4. Each block contained the same number of convolutional operations (see Fig. 3) and parameters configuration. All networks were trained in 3 anatomical views (axial, coronal, and sagittal) from scratch with the same training data-scheme; each data-configuration was carried out four different times, and the best weights were selected based on performance in the validation set. Furthermore, the 2D models were implemented with the same multi-slice input method, and 3D models were trained in different anatomical views by permuting the axis from the input volumes just like their 2D counterparts. Finally, all comparative models were implemented with the above-mentioned ensemble approach (see Section 2.2.3), and segmentation performance on the unseen test set was evaluated by comput-

Table 2

Mean (and standard deviation) of segmentation performance metrics of the F-CNN models. Models were ranked ascendingly by individual metrics and the overall rank (geometric mean of the metric rankings). We show significance indicators of the paired Wilcoxon signed-rank test comparing the proposed *AttFastSurferCNN* vs. benchmarked F-CNNs. Note *FastSurferCNN* is abbreviated to FSCNN and *AttFastSurferCNN* to AttFSCNN.

	Dice		VS		AVD (mm)		Overall Rank
	Mean (SD)	Rank	Mean (SD)	Rank	Mean (SD)	Rank	
AttFSCNN	0.8525 0.0561	6	0.9104 0.0634	6	0.2154 0.1530	5	5.65
FSCNN	0.8506 0.0577	5*	0.9081 0.0658	4	0.2134 0.1488	6	4.93
QuickNat	0.8506 0.0555	5*	0.9084 0.0635	5*	0.2174 0.1469	4*	4.64
UNet	0.8473 0.0610	3**	0.9071 0.0670	3*	0.2218 0.1567	3**	3.00
FSCNN3D	0.8163 0.0944	2**	0.8794 0.1109	1**	0.2510 0.1821	2**	1.59
UNet3D	0.8038 0.0820	1**	0.8878 0.0950	2**	0.2549 0.1582	1**	1.26

Significance: * $p < 0.05$, ** $p < 0.01$

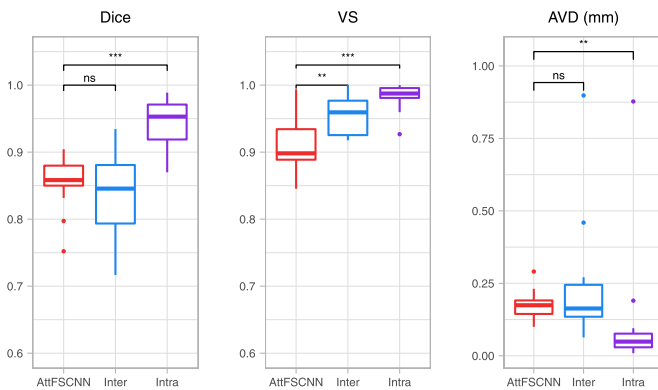


Fig. 7. Segmentation similarity scores for total OB comparing *AttFastSurferCNN* (AttFSCNN) vs. manual raters (intra- and inter-rater scores), as well as significance level indicators (paired two-sided Wilcoxon signed-rank). Significance: *** $p < 0.001$, ** $p < 0.01$, ns : $p \geq 0.05$.

ing three similarity metrics (Dice, AVD, and VS) between the predicted maps and manuals labels.

In Table 2 we present the similarity scores for OB tissue segmentation of all evaluation metrics as well as individual and overall ascending rankings and significance indicators of the two-sided Wilcoxon signed-rank test comparing the proposed *AttFastSurferCNN* vs. benchmarked F-CNNs. Here, we observed that our proposed *AttFastSurferCNN* has the highest overall ranking. Additionally, *AttFastSurferCNN* outperforms all other benchmark networks in all comparative metrics with statistical significance ($p < .05$) except for *FastSurferCNN*. *FastSurferCNN* outranks our proposed method in AVD, however, there is no statistical difference between them. On the other hand, *AttFastSurferCNN* outperforms *FastSurferCNN* in Dice and VS with a statistical significance ($p < .05$) in Dice. Finally, it is important to note that all 2D approaches drastically outperform the 3D models with up to 3% improvement of the Dice, 2.7% of VS and 4% of AVD between *UNet* (the lowest rank 2D model) and *3D-FastSurferCNN* (the highest rank 3D model).

Finally, to put the *AttFastSurferCNN* results into context, we compared the performance against the inter and intra-rater variability scores obtained in the manual annotation reproducibility experiment. For a fair comparison, this analysis is exclusively done in 19 cases that are also part of the test-set. Fig. 7 presents box plots for the three accuracy metrics as well as statistical significance indicators (paired two-sided Wilcoxon signed-rank test). We observed that *AttFastSurferCNN* results are significantly lower than the intra-rater scores. However, this was expected as we used the main-rater labels to train our F-CNNs and the

intra-rater scores are usually very difficult to reach for an automated method. Moreover, the proposed network outperforms the inter-rater scores (Dice: **0.8566** vs. 0.8386, and AVD: **0.1745 mm** vs. 0.2264 mm) in localizing the OB tissue and recognizing its boundaries, even if no statistical significance can be inferred from the statistical test. On the other hand, for VS, the inter-rater results are significantly better (VS: 0.9115 vs. **0.9555**); nevertheless, there is an average VS difference of only 0.04 between label maps translating to a small volume discrepancy of around 0.020 mm³ by every segmented voxel.

3.2.3. Ensemble

In this section, we tested our ensemble approach of combining the output of four *AttFastSurferCNN* against each individual *AttFastSurferCNN* trained in the previous section. We observed that all standalone models have comparable results in the three similarity metrics (Dice, VS, and AVD) as shown in Table 3. Thus OB segmentation knowledge is not driven by any particular data-subset, and all *AttFastSurferCNN*s outperform the inter-rater scores for Dice (0.8386) and AVD (0.2264 mm). Furthermore, the proposed ensemble model significantly outperforms all standalone (non-ensembled) models with respect to Dice and AVD ($p < .05$, paired two-sided Wilcoxon signed-rank test). We observed no statistical difference between models in VS except for *AttFastSurferCNN-4* where the proposed merged method has better results. Finally, we empirically observed that the ensemble model smoothes the label maps slightly, resulting in visually more appealing boundaries as illustrated in Fig. 8.

3.3. Age and sex effects sensitivity (E3)

OB volumes obtained from manual segmentations of T2w images have shown to be negatively correlated with age (Buschhüter et al. (2008); Hummel et al. (2011, 2015)). Therefore, any automated method that intends to detect this small structure should be able to replicate these effects. As a consequence, we evaluated the sensitivity of our proposed pipeline to replicate ground truth age dependencies in the in-house unseen test-set ($n = 203$) which has a comparable size to other manually annotated OB datasets (Buschhüter et al. (2008); Hummel et al. (2015)) used for volume correlations. Furthermore, we compared our results with the F-CNNs used in the benchmark (see Section 3.2.2). The association of OB volumes (OBV) and age was assessed using a linear regression after accounting for sex and head-size (estimated total intracranial volume, eTIV) (Model: $OBV \sim age + sex + eTIV$). All statistical analyses were performed in R (R Core Team (2020)) and eTIV estimations were computed using *FreeSurfer* (Buckner et al. (2004); Fischl (2012); Fischl et al. (2002)).

Table 3

Mean (and standard deviation) of segmentation performance metrics of the proposed ensemble approach and single *AttFastSurferCNN* (AttFSCNN) models. Models were ranked ascendingly by individual metrics and the overall rank (geometric mean of the metric rankings). We show significance indicators of the paired Wilcoxon signed-rank test comparing the proposed ensemble *AttFastSurferCNN* vs. single *AttFastSurferCNN*.

Model	Dice		VS		AVD (mm)		Overall Rank
	Mean(SD)	Rank	Mean(SD)	Rank	Mean(SD)	Rank	
Ensemble	0.8525	5	0.9104	3	0.2154	5	4.22
AttFSCNN	0.0561		0.0634		0.1530		
AttFSCNN	0.8482	4**	0.9112	4	0.2225	4*	4.00
3	0.0589		0.0659		0.1706		
AttFSCNN	0.8477	3**	0.9096	2	0.2234	2**	2.29
2	0.0578		0.0646		0.1614		
AttFSCNN	0.8476	2**	0.9115	5	0.2276	1**	2.15
1	0.0552		0.0625		0.1749		
AttFSCNN	0.8469	1**	0.9077	1*	0.2230	3**	1.44
4	0.0580		0.0666		0.1491		

Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

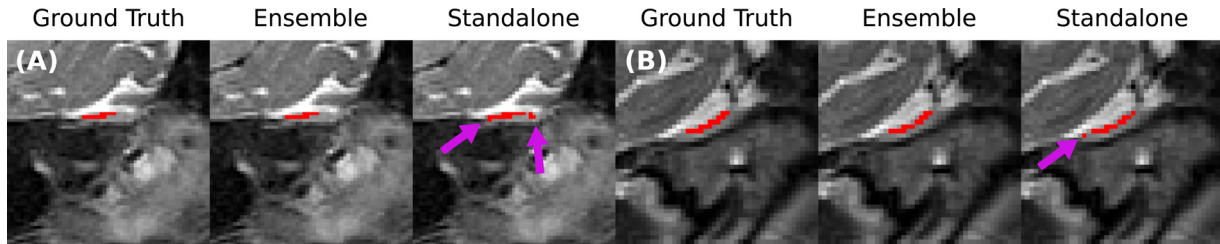


Fig. 8. Comparison of the manual ground truth vs. predictions of the right OB from two subjects on sagittal T2-weighted MRI of the in-house test-set. Purple arrows indicate where the proposed ensemble *AttFastSurferCNN* improves the segmentation over a standalone *AttFastSurferCNN*.

Table 4

Association of OB volumes (OBV) and age after accounting for sex and head-size (eTIV) on the in-house test-set for the manual labels (ground truth) and benchmark networks. Linear regression model: $OBV \sim age + sex + eTIV$. Note *FastSurferCNN* is abbreviated to FSCNN and *AttFastSurferCNN* to AttFSCNN.

	Ground Truth	AttFSCNN	FSCNN	QuickNat	UNet	FSCNN3D	UNet3D
(Intercept)	53.292** (1.953)	55.517*** (1.636)	54.774*** (1.620)	56.038*** (1.642)	55.330*** (1.638)	45.714*** (1.535)	47.186*** (1.501)
Age	-0.319*** (0.092)	-0.232** (0.077)	-0.204** (0.076)	-0.213** (0.077)	-0.211** (0.077)	-0.225** (0.072)	-0.241*** (0.070)
Sex:	5.940 (3.463)	3.150 (2.900)	2.612 (2.871)	3.409 (2.910)	3.017 (2.903)	1.980 (2.721)	2.897 (2.660)
m/f	14.286 (10.238)	32.189*** (8.577)	32.297*** (8.490)	31.713*** (8.605)	32.590*** (8.586)	25.022** (8.047)	21.116** (7.867)
eTIV	0.124	0.205	0.193	0.199	0.199	0.157	0.156
R-squared	0.124	0.205	0.193	0.199	0.199	0.157	0.156
N	203	203	203	203	0.203	0.203	203

Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

All predicted OB volumes significantly decreased with age as can be seen in Table 4, which in turn follows the behavior of the manual data and other studies (Buschhüter et al. (2008); Hummel et al. (2011, 2015)). We found an improvement in the modeling (R^2) of the age effects in the *AttFastSurferCNN* compared to the ground truth and the other comparative networks. Finally, we did not find a sex difference for any of the models, and, as expected, the inferred OBV are positively associated with eTIV (see Table 4).

3.4. E4: No apparent olfactory bulb (E4)

As the proposed pipeline is to be deployed as a post-processing OB analysis pipeline for the T2w MRI of the Rhineland Study, it should be robust to cases without an apparent OB that - based on the size of our in-house dataset - occur with an approximate prevalence of 2%. In this section, we processed the 12 flagged cases with no apparent OB and evaluated the OB volume estimates. Note, all cases used for training our *AttFastSurferCNN* have a visible OB.

The automated method agreed with the main-rater in 50% percent of these cases as illustrated in Fig. 9 B) and shown in Appendix Fig. A2. For the remaining cases: three had a total predicted volume smaller than

2.5 mm^3 and the other three between 7 mm^3 to 10.2 mm^3 . We additionally observed that there is hemisphere asymmetry where the maximum predicted volume by any hemisphere was 8.7 mm^3 translating in a detection of only 17 voxels. After visually inspecting the predicted label maps by two different raters, we observed that with the current resolution our raters cannot reliably assess the predicted segmentation of an individual olfactory bulb with a size smaller than 10 mm^3 as seen in Fig. 9 C) and D) where the in-plane segmentation is only a few voxels. For this reason, we additionally evaluated the effects of OB size on the segmentation accuracy of the automated method for the test-set. We observed that segmentation performance decreases in subjects with a total OB smaller than 20 mm^3 . Furthermore, OB volumes are positive correlated with similarity metrics (Dice: $R = 0.39, p < .001$, VS: $R = 0.23, p < .001$) and negative correlated with AVD ($R = -0.39, p < .001$), a dissimilarity metric. For more detailed information see Appendix Fig. A3.

3.5. Sequence stability (E5)

In this section, we processed all $T2w^a$ and $T2w^b$ scans from the stability dataset with the proposed pipeline. Afterwards, we assessed the pipeline stability by comparing the similarity of total OB volume across

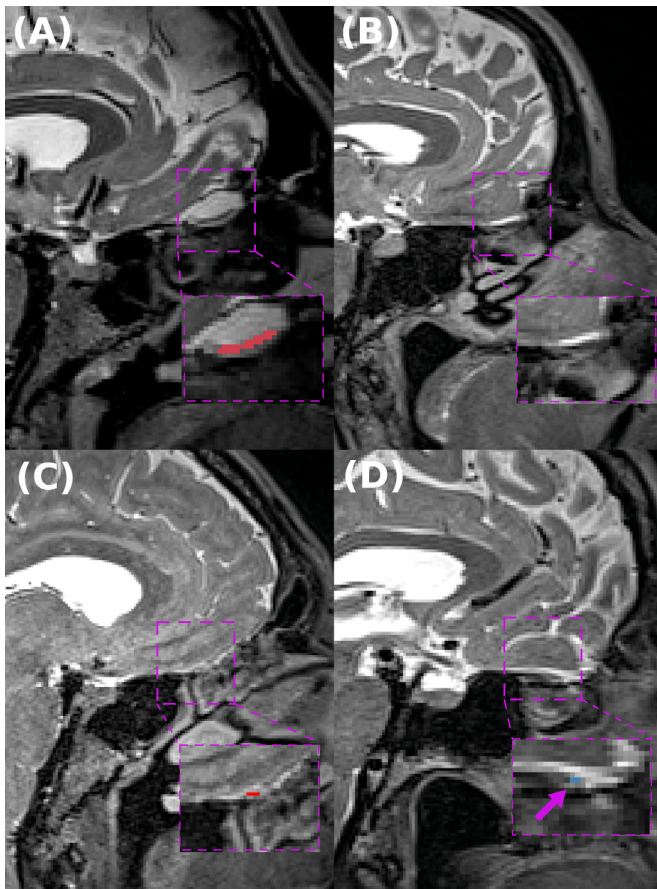


Fig. 9. A-D) Sagittal T2-weighted MR images and predictions on cases from the Rhineland Study. A) Normal subject from the in-house dataset with a visible OB, B) Subject without an apparent OB where the pipeline also agrees with our main rater. C-D) Subjects flagged with no visible OB by our main rater, however, the pipeline still predicts some voxels as OB (total volume < 10.2 mm³) due to the current resolution our raters cannot reliably assess the predicted segmentation. Note, red indicates Right OB and blue left OB (purple arrow indicate the segmented voxel). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

sequences by volume similarity (VS) as described in the metric evaluation section. Additionally, we calculated the agreement of total OB volume between sequences by an intra-class correlation (ICC) using a two-way fixed, absolute agreement and single measures with a 95% confidence interval (ICC(A,1) (McGraw and Wong (1996))). To further compare the agreement between sequences, three random subjects from the stability dataset were selected and both T2w sequences were manually annotated. Subsequently, segmentation performance metrics (Dice, VS, AVD) between the manual and predicted label maps were computed. It is important to note that we did not compute overlap segmentation performance metrics (Dice and AVD) across different sequence label maps of the same subject as this would require registering the scans. It would not only include inherent variance from acquisition noise (e.g. motion artefacts, non-linearities based on different positioning) but also variance due to registration inaccuracies and interpolation artefacts.

After visual quality inspection, a total of 7 scans were excluded from this analysis due to image artefacts such as motion or low contrast (see Appendix Fig. A4 for two examples). For the remaining cases (n=102), we observed a good agreement between the $T2w^a$ and $T2w^b$ sequences (ICC: 0.897 [0.845 - 0.931]) and a volume similarity (VS: 0.889 (0.090)) comparable to the one described in previous sections. However, we observed a statistical difference between volume estimates ($p < .01$, paired two-sided Wilcoxon signed-rank test). Furthermore, to give more context on how variations in a T2w sequence affect the pipeline's predic-

tions, we analyzed the segmentation similarity on the manually annotated subset. As expected, the result on the $T2w^a$ (training) sequence outperforms the $T2w^b$ segmentation results (Dice: **0.8622** vs. 0.8597, VS: **0.9343** vs. 0.9066 and AVD: **0.1816** mm vs. 0.1965 mm). Nevertheless, the segmentation performance in both sequences is in the range of intra-rater scores (Dice: 0.8386, VS: 0.9555, and AVD: 0.2264 mm). Demonstrating that systematic sequence improvements can be beneficial in an ongoing population study without diminishing the performance of the proposed method. Even though our pipeline showed volume stability across sequences and that segmentation performance was not affected, it is still important to control for MRI sequence in any downstream statistical analysis when including data from multiple MRI sequences.

3.6. Generalizability (E6)

The lack of MR hardware heterogeneity (i.e. scanners, field strength, and acquisition parameters) in our training set can limit the ability of the neural network to generalize to unseen T2w images acquired under different conditions. In order to quantify the robustness of our pipeline, we tested it on 30 subjects of the HCP dataset, acquired with a different resolution with isotropic 0.7 mm voxels. In addition to sequence differences, HCP images are de-faced. In order to analyze our method at the native 0.7 mm HCP resolution as well as at the default 0.8 mm network resolution, we constructed manual annotations twice per subject, one for each resolution.

We perform three experiments: A) Input images were resampled to the default network resolution (isotropic 0.8 mm), resulting label maps were upsampled to the original 0.7 mm resolution and compared to the manual reference there. B) Images were processed directly at the native resolution of 0.7 mm and compared to the 0.7 mm manual reference, thus, evaluating the networks' generalizability to segment inputs at a slightly higher and unseen resolution directly. C) Same as A) but instead of upsampling final labels they are compared with the manual reference delineated at 0.8 mm, avoiding the final upsampling step. This permits quantifying the accuracy for the default behaviour of the network, if final segmentations at 0.8 mm are sufficient for the user.

Fig. 10 clearly indicates that option A (orange) provides the lowest performance, most likely due to the fact that it includes interpolation artefacts from upsampling the final labels. Resampling label maps is often problematic and should be avoided. If final results are required at the original (here 0.7 mm) resolution it is indeed better to directly segment these images at the native resolution (option B, blue boxes). Even though the network has not been trained on this resolution, it can generalize remarkably well. Option C demonstrates that best results can be obtained at the default network resolution of 0.8 mm, which is the recommended approach.

As expected, overall performance on HCP data is slightly lower than the results obtained on our in-house dataset (see Section 3.2.2). The HCP dataset, however, consists of de-faced scans (never encountered during training) from a younger age distribution, and was acquired with different acquisition parameters. Due to these differences, segmentation scores are not directly comparable. Nevertheless, the proposed pipeline generalizes quite well across age-categories, especially when evaluated at the original training resolution as metrics remained relatively stable with an overall good performance (Dice: 0.7816, VS: 0.8583, and AVD: 0.2683 mm, red boxes). Additionally, we observe that segmentation accuracy decreases slightly for ages outside the training range (namely 22 to 25, training data started at age 30). Yet the overall high accuracy shows that our proposed pipeline can robustly generalize to the unseen HCP data. Examples of OB segmentations for both the in-house as well as the HCP dataset can be found in Fig. 11.

4. Discussion

In this work, we established, validated, and implemented a novel deep learning pipeline to segment and quantify the olfactory bulb on

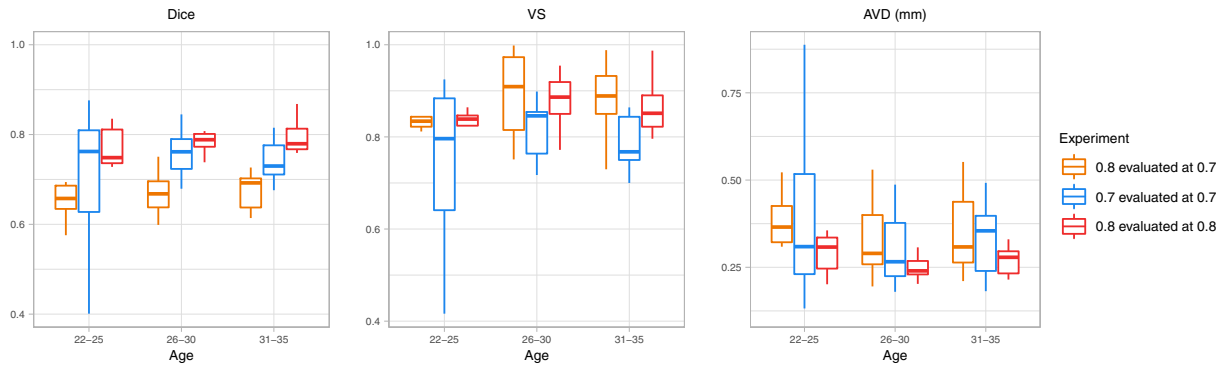


Fig. 10. Segmentation similarity scores of total OB for the 30 labelled cases from the HCP dataset stratified by age category, as well as comparison of the pipeline's performance at native HCP resolution (0.7 mm isotropic, with upsampling: orange, directly: blue) and at the networks original training resolution (0.8 mm isotropic, red). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

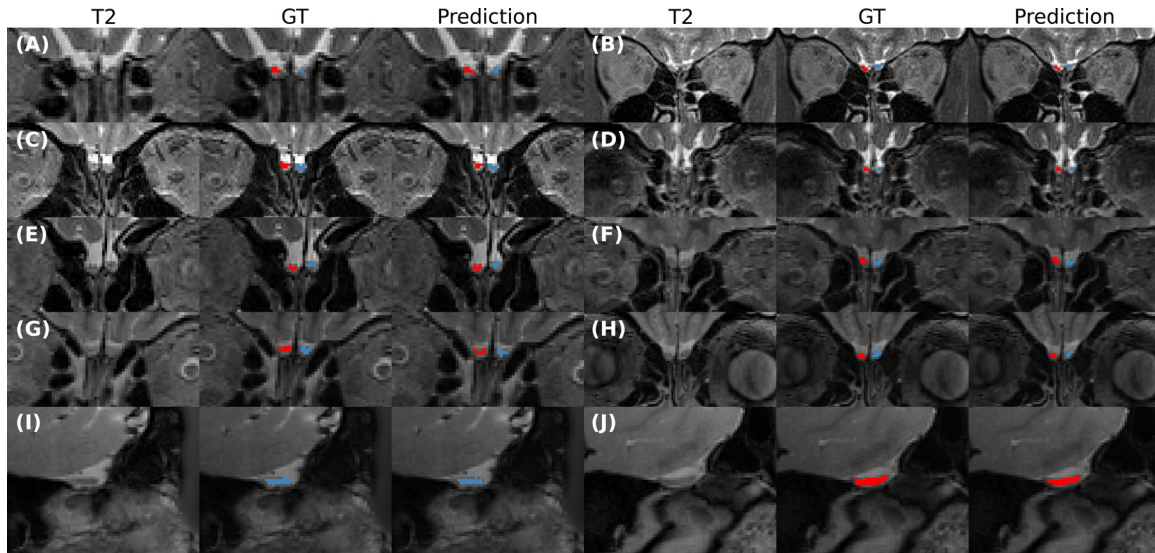


Fig. 11. Comparison of the ground truth vs. predictions on coronal (A-H) and sagittal (I-J) T2w MRI from subjects of the Rhineland Study (A-E) and HCP (F-J) dataset at 0.8 mm. A-J) Accurate automatic segmentation of total OB on a heterogeneous population. Note, blue: left OB and red: right OB. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

high resolutional T2-weighted MR scans. The proposed pipeline is fully automatic and can analyze a 3D volume in less than a minute in an end-to-end fashion, even though it implements a three-stage design. The use of deep learning components for localizing and segmenting the OB enables the pipeline to accurately and quickly quantify the OB volume, providing a robust and reliable solution for assessing OB volumes in a large cohort study.

Segmenting the OB in T2w scans is a challenging task due to size, sensitivity to artefacts, age effects, and visibility on MR images (partial volume effects). Despite all these challenges, we demonstrate the feasibility of segmenting the OB on high resolutional isotropic T2w MR images. Our main rater's manual annotations exhibit a high intra-rater reliability in terms of boundary delineation, OB localization, and volume estimation. Furthermore, we verified the reproducibility of our labeling protocol with inter-rater reliability similar to the one reported in other manually annotated medical datasets (Billot et al. (2020); Estrada et al. (2020)). We cannot directly compare the segmentation performance with other studies that manually labeled the OB on T2w MR images as they only report the volume difference for repeated measurements by a single observer or across observers (Hummel et al. (2011); Joshi et al. (2020); Mueller et al. (2005); Yousem et al. (1997)). Nonetheless, the volume similarity for both inter and intra-rater variability yields comparable or even better results than

the OB studies mentioned above. These results demonstrate the quality of the manual annotations and soundness of developing an automated method for segmenting the OB using a supervised learning technique.

For the first stage of the pipeline, i.e. localization of the OB in a whole-brain image, all four implemented *FastSurferCNN*s can successfully localize a forebrain region containing the OBs from both hemispheres (region of interest) and determine a cropping coordinate based on the centroid from a segmentation prediction map. However, for our final localization model, we chose the *FastSurferCNN* model 4 as it out-ranked all other models in all evaluation metrics (Dice, VS, and R). The implemented localization block is able to identify the region of interest in a low-resolution image even when the input scans are defaced as in the HCP dataset or have motion artefacts as illustrated in Appendix Fig. A4.

For the more challenging task of segmenting OB, we contribute a deep learning architecture (*AttFastSurferCNN*) by incorporating a self-attention module inside our *FastSurferCNN*. The introduction of a self-attention mechanism improves the network's modeling of global dependencies (Fu et al. (2019); Zhang et al. (2019)), thus increasing the attention to spatial information and boosting the learning of such a fine-grained structure as the OB. We demonstrate that *AttFastSurferCNN* recovers OB significantly better than the standard *FastSurferCNN* and other traditional deep learning variants used for semantic segmentation. It is

also important to note that our proposed method shows an improvement when evaluating volume associations in a large cohort despite the slight changes at the image metric level. Additionally, each of the four individual *AttFastSurferCNNs* that composes the ensemble model outperforms manual inter-rater scores for segmenting and delineating the OB. Even though the volume similarity from the proposed method is lower than the one from the manual raters, the mean volume difference ($\approx 9\%$) is still in the 10% acceptable difference used as selection criteria by other studies for including the OB volumes of a subject with multiple manual annotations (Joshi et al. (2020)). Moreover, the implemented ensemble approach regularizes the predicted segmentation by combining the spatial context from different views and models, ultimately improving the segmentation of the OB boundaries and reducing the variance due to networks initialization. Furthermore, the predicted probability maps from all individual *AttFastSurferCNNs* can be used to compute the pipeline uncertainty (Kendall and Gal (2017); Roy et al. (2019a)), a potential quality control marker for flagging problematic cases.

The 2.5D approach used for all 2D benchmark networks of multi-network view-aggregation and multi-slice input drastically outperforms the comparative 3D models. Showing that 3D methods are not always the best method and that 2D models can yield better results when strategies to increase the spatial information are included as the one used in this work. Moreover, reducing the scope of the local neighbourhood when segmenting a small structure like the OB is beneficial as it reduces the amount of redundant information and increases the attention to the spatial information surrounding the OB. To improve attention in a 3D network towards OB, a naive solution would be to include the proposed self-attention layer. However, the computation of an attention map of size $N \times N$, where N are the number of voxels, will considerably increase the GPU memory requirements and 3D networks are inherently memory expensive to train. Therefore, a self-attention layer is not an efficient and scalable solution for this type of networks. More efficient techniques are being studied, but they are outside the scope of this paper.

As demonstrated in the Rhineland data, the proposed pipeline successfully identifies the OB on a T2w scan as seen in Fig. 11 A) to E). The pipeline also replicates the negative correlation of OB volumes with age reported in previous studies (Buschhüter et al. (2008); Hummel et al. (2011, 2015)) and also visible in our manual annotations. We, furthermore, detected no sex difference after accounting for head size, however, estimates from *AttFastSurferCNN* and all comparative networks are positively correlated with head size - a result that is also detected in the manual segmentations - as expected - but with a lower significance and magnitude. All automated methods show stronger and less variable eTIV effect across subjects (see Table 4), explaining the significance discrepancy. The difference in effect magnitudes can be attributed to the FCNN's ability to learn consistent information across subjects exhibiting stability to random noise and thus generating smoother segmentations than manual raters. Furthermore, our proposed pipeline efficiently handles cases without an apparent OB by not segmenting the structure at all or only a few voxels ($< 10 \text{ mm}^3$) as seen in Fig. 9 B), C), and D). Additionally, the sequence stability dataset demonstrates a good agreement of volume estimates between sequences. It must be noted that the difference in volume estimates includes not only potential variances of the processing pipelines but also variance from acquisition noise (e.g. motion artefacts, non-linearities based on different head positions). Therefore we recommend controlling for MRI sequence in follow-up statistical analysis when pooling input data. As consistent changes in a sequence can reflect a consistent change in measured OB size. Nonetheless, segmentation performance in all sequences yields comparable results to the manual inter-rater scores. The fact that our results in the Rhineland Study data (i) replicate known OB volume effects, (ii) properly identify scans without an apparent OB, and (iii) demonstrate a good agreement of volume estimates among variations of the study's T2w sequence corroborates robustness and stability of our pipeline. Nevertheless, due to current image resolution and based on quality assessment of all the predicted label maps generated in this work, we recommend visually in-

spect cases with an OB volume smaller than 20 mm^3 before including them in any analysis.

Our automated method not only exhibits generalizability across a wide range of ages from the Rhineland Study but can also extend to another population distribution (HCP dataset) with different acquisition parameters. The pipeline presents optimal results when the input images have the default training resolution of 0.8 mm isotropic. Nonetheless, results at a different resolution (HCP native resolution of 0.7 mm) still yield a good performance even with all the various other differences, e.g. different distribution, de-faced image, acquisition parameters, and image resolution. Even though our method shows robustness to de-facing pre-processing steps in HCP, de-facing or skull stripping can be problematic due to the proximity of the OB region to the cropped mask, in the worst case scenario - depending on the method - resulting in accidentally cropping into the OB. Therefore, full head T2w scans are the recommended input to our pipeline. Additionally, T2w scans with a different resolution from the ones presented in this work can also be analyzed by running the pipeline with the default behaviour (resampling inputs to 0.8 mm) or by processing inputs directly at the native image resolution, if it is close to 0.8 mm isotropic. In these cases is highly recommended, however, that segmentation quality is assessed by the user. Generally, since the pipeline is based on deep learning, the model can easily be fine-tuned to another desired resolution by retraining or by more aggressive scaling augmentation techniques.

In conclusion, we have developed a fully automated post-processing pipeline for OB segmentation on sub-millimeter T2-weighted MRI based on advanced deep learning methods. To the best of our knowledge, the presented pipeline is the first to accurately segment the OB in a large cohort and is meticulously validated not only against segmentation accuracy but also with respect to known OB volume effects (e.g. age).

Credit authorship contribution statement

Santiago Estrada: Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Ran Lu:** Investigation, Data curation, Writing – original draft, Writing – review & editing. **Kersten Diers:** Validation, Formal analysis, Writing – review & editing. **Weiyi Zeng:** Data curation, Writing – review & editing. **Philipp Ehse:** Data curation, Writing – original draft. **Tony Stöcker:** Data curation, Resources, Writing – review & editing. **Monique M. B. Breteler:** Conceptualization, Supervision, Funding acquisition, Resources, Writing – review & editing. **Martin Reuter:** Conceptualization, Validation, Resources, Writing – original draft, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Acknowledgment

We would like to thank the Rhineland Study group for supporting the data acquisition and management. This work was supported by DZNE institutional funds, the Federal Ministry of Education and Research (BMBF), Germany (grant numbers: 031L0206 and 01GQ1801), the Diet-Body-Brain Competence Cluster Nutrition Research funded by the BMBF (grant numbers: 01EA1410C and 01EA1809C), and the NIH (grant numbers: R01NS083534, R01LM012719, P41EB01589620). Ran Lu was partially supported by a scholarship from China Scholarship Council.

External data used in the preparation of this work were obtained in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the [NIH Blueprint for Neuroscience Research](#); and by the McDonnell Center for Systems Neuroscience at Washington University.

Appendix A

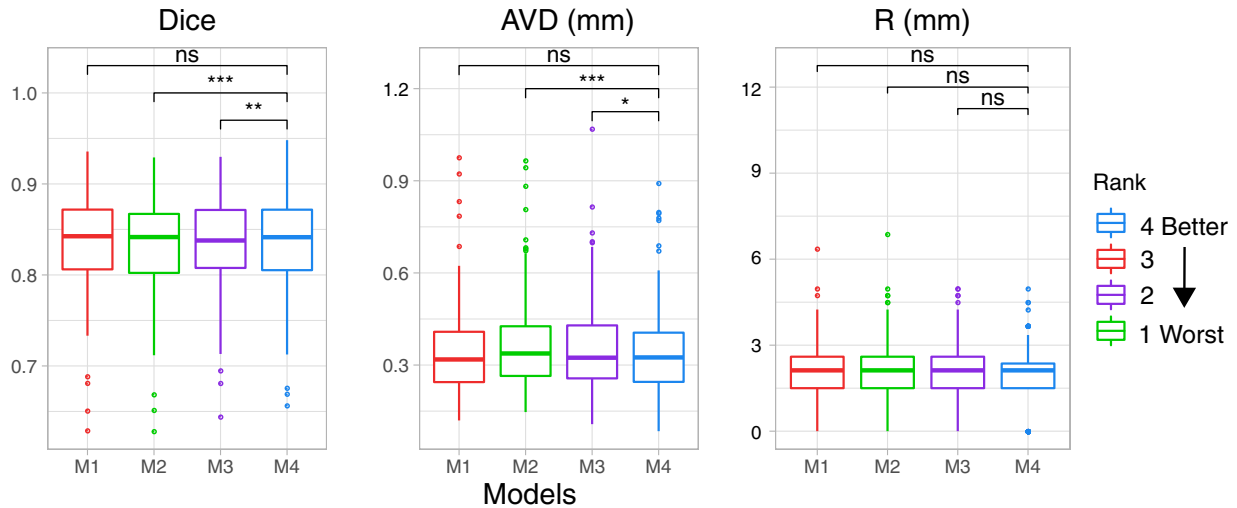


Fig. A1. Similarity metrics scores for ROI localization comparing all trained *FastSurferCNN* models. Models were ranked ascendingly by individual metrics (box-plot color) and the overall rank (geometric mean of the metric rankings). We show significance level indicators of the paired Wilcoxon signed-rank test comparing *FastSurferCNN*-4 (M4, model with best overall rank) against the other *FastSurferCNN*s (M1,M2,M3). Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, ns : $p \geq 0.05$.

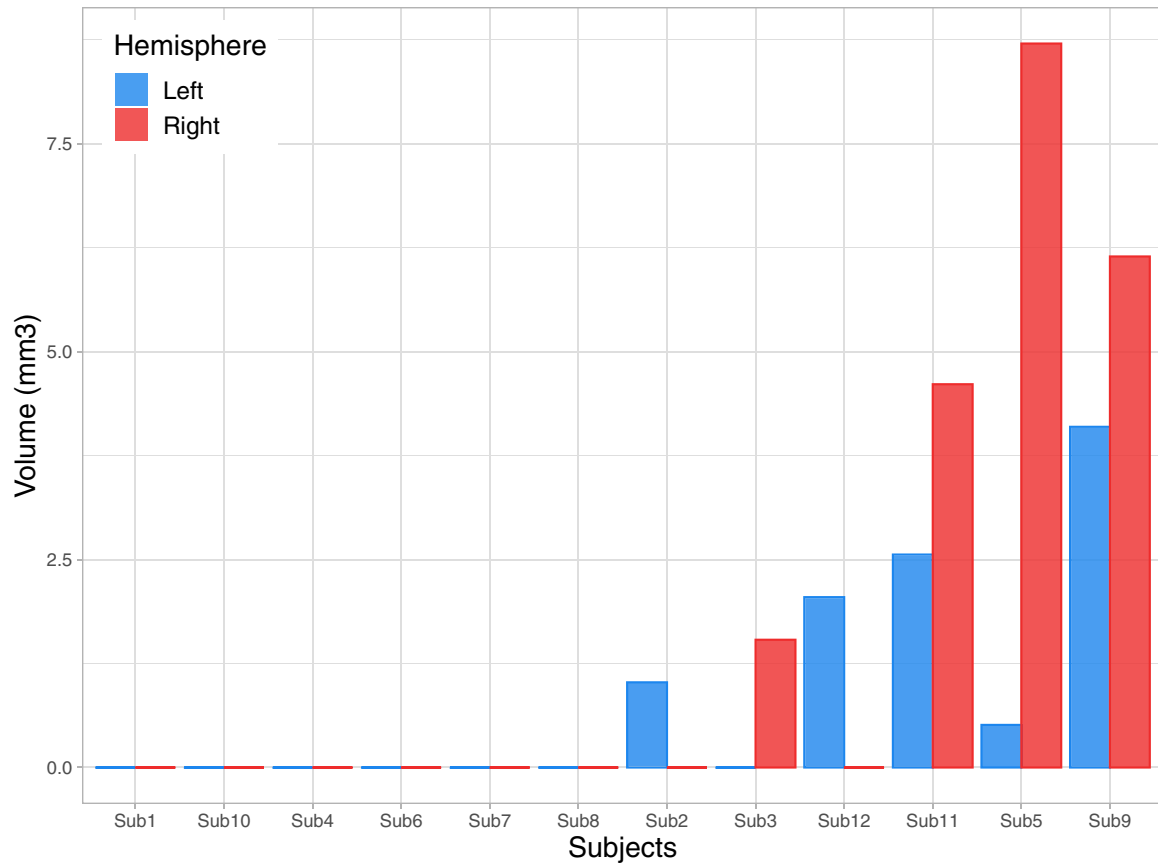


Fig. A2. OB volume estimates (left (blue) and right (red)) after processing the 12 subjects flagged with no visible OB with the proposed automated pipeline. The automated method agreed with the main-rater in 50% percent of these cases. For the remaining cases: three had a total predicted volume smaller than 2.5 mm³ and the other three between 7 mm³ to 10.2 mm³. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table A1
OB demographics for the total in-house dataset and for the training and testing subsets. Descriptive data were expressed as mean (SD) or count (percentage) for continuous or categorical variables, respectively. Inter group differences were compared with the Student's *t*-test for continuous variables and with the Pearson's chi-square test for categorical variables.

	Trainset				Testset (N=203)	Total (N=560)	p value
	Split_1 (N=90)	Split_2 (N=89)	Split_3 (N=89)	Split_4 (N=89)			
Sex							0.996
Female	52 (57.8%)	52 (58.4%)	50 (56.2%)	50 (56.2%)	114 (56.2%)	318 (56.8%)	
Male	38 (42.2%)	37 (41.6%)	39 (43.8%)	39 (43.8%)	89 (43.8%)	242 (43.2%)	0.992
Age							
Mean (SD)	53.900 (12.986)	53.360 (13.487)	53.708 (13.345)	54.348 (12.763)	53.837 (13.540)	53.832 (13.247)	0.126
Range	30.000 - 81.000	31.000 - 85.000	30.000 - 82.000	31.000 - 83.000	30.000 - 87.000	30.000 - 87.000	
OB Volume(mm^3)							
Mean (SD)	52.173 (14.623)	53.064 (15.814)	55.342 (13.353)	51.424 (13.629)	55.896 (18.576)	54.049 (16.085)	
Range	19.456 - 84.480	24.576 - 88.064	29.696 - 84.992	21.504 - 84.480	12.800 - 111.104	12.800 - 111.104	

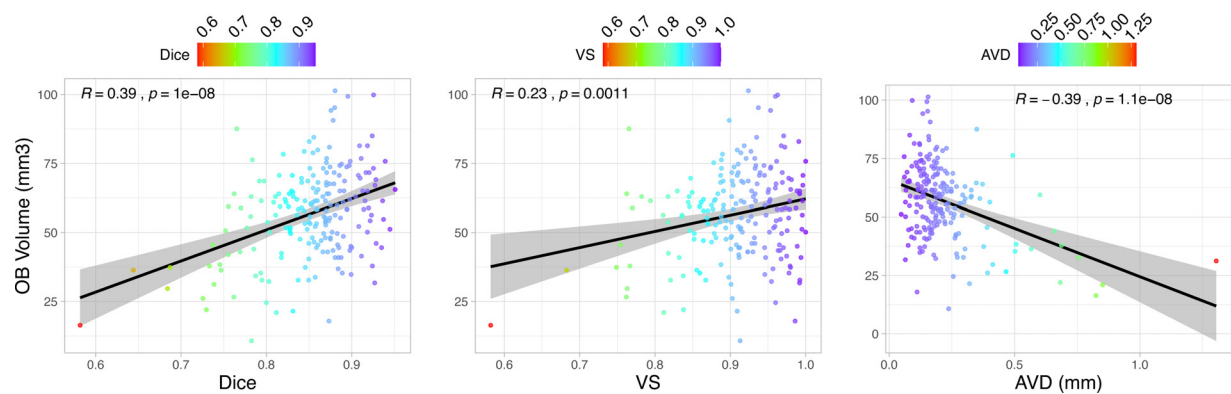


Fig. A3. Scatterplots of OB volume estimates and segmentation similarity metrics on the in-house test-set as well as the Pearson correlation coefficient and linear regression. We observed that segmentation performance decreased with OB size. Especially in subjects with a total OB volume smaller than 20 mm³.

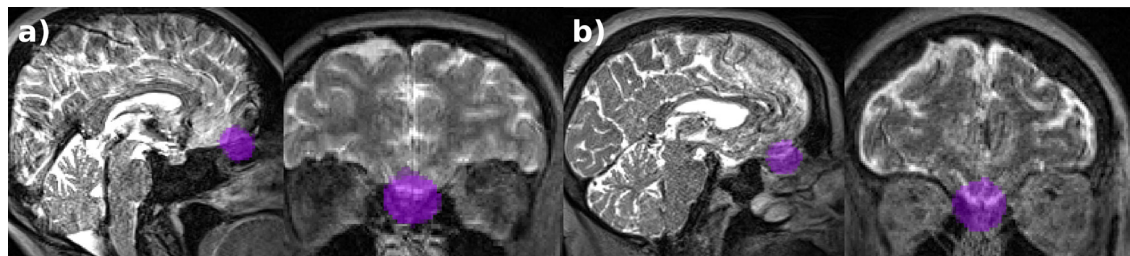


Fig. A4. Sagittal and Coronal T2-weighted MR images and predictions from the localization stage (purple) on two cases from the Rhineland Study. A-B) Present subjects excluded from the volume estimates sequence stability analysis (E5) due to severe motion artefact. Nonetheless, the localization stage still can detect a region containing both OBs. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

References

- Attems, J., Walker, L., Jellinger, K.A., 2014. Olfactory bulb involvement in neurodegenerative diseases. *Acta Neuropathol.* 127 (4), 459–475.
- Avants, B.B., Tustison, N., Song, G., 2009. Advanced normalization tools (ants). *Insight j* 2 (365), 1–35.
- Badrinarayanan, V., Kendall, A., Cipolla, R., 2017. Segnet: a deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans Pattern Anal Mach Intell* 39 (12), 2481–2495.
- Billot, B., Bocchetta, M., Todd, E., Dalca, A.V., Rohrer, J.D., Iglesias, J.E., 2020. Automated segmentation of the hypothalamus and associated subunits in brain mri. *Neuroimage* 223, 117287.
- Breteler, M.M., Stöcker, T., Pracht, E., Brenner, D., Stirnberg, R., 2014. Mri in the rhineland study: a novel protocol for population neuroimaging. *Alzheimer's & Dementia: The Journal of the Alzheimer's Association* 10 (4), P92.
- Buckner, R.L., Head, D., Parker, J., Fotenos, A.F., Marcus, D., Morris, J.C., Snyder, A.Z., 2004. A unified approach for morphometric and functional data analysis in young, old, and demented adults using automated atlas-based head size normalization: reliability and validation against manual measurement of total intracranial volume. *Neuroimage* 23 (2), 724–738.
- Buschhüter, D., Smitka, M., Puschmann, S., Gerber, J.C., Witt, M., Abolmaali, N., Hummel, T., 2008. Correlation between olfactory bulb volume and olfactory function. *Neuroimage* 42 (2), 498–502.
- Busse, R.F., Brau, A.C., Vu, A., Michelich, C.R., Bayram, E., Kijowski, R., Reeder, S.B., Rowley, H.A., 2008. Effects of refocusing flip angle modulation and view ordering in 3d fast spin echo. *Magn Reson Med* 60 (3), 640–649.
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O., 2016. 3d u-net: learning dense volumetric segmentation from sparse annotation. In: *International conference on medical image computing and computer-assisted intervention*. Springer, pp. 424–432.
- Croy, I., Nordin, S., Hummel, T., 2014. Olfactory disorders and quality of life-An updated review. *Chem. Senses* 39 (3), 185–194.
- Dintica, C.S., Marseglia, A., Rizzuto, D., Wang, R., Seubert, J., Arfanakis, K., Bennett, D.A., Xu, W., 2019. Impaired olfaction is associated with cognitive decline and neurodegeneration in the brain. *Neurology* 92 (7), e700–e709.
- Dong, H., Yang, G., Liu, F., Mo, Y., Guo, Y., 2017. Automatic brain tumor detection and segmentation using u-net based fully convolutional networks. In: *annual conference on medical image understanding and analysis*. Springer, pp. 506–517.
- Doty, R.L., 2017. Olfactory dysfunction in neurodegenerative diseases: is there a common pathological substrate? *The Lancet Neurology* 16 (6), 478–488.
- Estrada, S., Conjeti, S., Ahmad, M., Navab, N., Reuter, M., 2018. Competition vs. concatenation in skip connections of fully convolutional networks. In: *International Workshop on Machine Learning in Medical Imaging*. Springer, pp. 214–222.
- Estrada, S., Lu, R., Conjeti, S., Orozco-Ruiz, X., Panos-Willuhn, J., Breteler, M.M., Reuter, M., 2020. Fatsegnet: a fully automated deep learning pipeline for adipose tissue segmentation on abdominal dixon mri. *Magn Reson Med* 83 (4), 1471–1483.
- Fischl, B., 2012. Freesurfer. *Neuroimage* 62 (2), 774–781.
- Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., Van Der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., et al., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33 (3), 341–355.
- Friston, K.J., 2003. Statistical Parametric Mapping. In: *Neuroscience databases*. Springer, pp. 237–250.
- Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H., 2019. Dual attention network for scene segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3146–3154.
- Goodfellow, I., Warde-Farley, D., Mirza, M., Courville, A., Bengio, Y., 2013. Maxout networks. In: *International conference on machine learning*. PMLR, pp. 1319–1327.
- Griswold, M.A., Jakob, P.M., Heidemann, R.M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., Haase, A., 2002. Generalized autocalibrating partially parallel acquisitions (grappa). *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 47 (6), 1202–1210.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- Henschel, L., Conjeti, S., Estrada, S., Diers, K., Fischl, B., Reuter, M., 2020. Fastsurfer-a fast and accurate deep learning based neuroimaging pipeline. *Neuroimage* 219, 117012.
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708.
- Hummel, T., Smitka, M., Puschmann, S., Gerber, J., Schaal, B., Buschhüter, D., 2011. Correlation between olfactory bulb volume and olfactory function in children and adolescents. *Exp Brain Res* 214 (2), 285.
- Hummel, T., Urbig, A., Huat, C., Duprez, T., Rombaux, P., 2015. Volume of olfactory bulb and depth of olfactory sulcus in 378 consecutive patients with olfactory loss. *J. Neurol.* 262 (4), 1046–1051.
- Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M., 2012. Fsl. *Neuroimage* 62 (2), 782–790.
- Joshi, A., Thaploo, D., Yan, X., Herrmann, T., Hummel, T., 2020. A novel technique for olfactory bulb measurements. *PLoS ONE* 15 (12), e0243941.
- Ju, C., Bibaut, A., van der Laan, M., 2018. The relative performance of ensemble methods with deep convolutional neural networks for image classification. *J Appl Stat* 45 (15), 2800–2818.
- Kamnitsas, K., Bai, W., Ferrante, E., McDonagh, S., Sinclair, M., Pawlowski, N., Rajchl, M., Lee, M., Kainz, B., Rueckert, D., et al., 2017. Ensembles of multiple models and architectures for robust brain tumour segmentation. In: *International MICCAI brainlesion workshop*. Springer, pp. 450–462.
- Kamnitsas, K., Ledig, C., Newcombe, V.F., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B., 2017. Efficient multi-scale 3d cnn with fully connected crf for accurate brain lesion segmentation. *Med Image Anal* 36, 61–78.
- Kendall, A., Gal, Y., 2017. What uncertainties do we need in bayesian deep learning for computer vision? In: *NIPS*, pp. 5580–5590. <http://papers.nips.cc/paper/7141-what-uncertainties-do-we-need-in-bayesian-deep-learning-for-computer-vision>.
- Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization. In: *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*.
- Liao, Z., Carneiro, G., 2016. On the importance of normalisation layers in deep learning with piecewise linear activation units. In: *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, pp. 1–8.
- Liao, Z., Carneiro, G., 2017. A deep convolutional neural network module that promotes competition of multiple-size filters. *Pattern Recognit* 71, 94–105.
- Lin, G., Shen, C., Van Den Hengel, A., Reid, I., 2016. Efficient piecewise training of deep structured models for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3194–3203.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2017. Focal loss for dense object detection. In: *Proceedings of the IEEE international conference on computer vision*, pp. 2980–2988.
- Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440. Boston, Massachusetts, USA.
- Loshchilov, I., Hutter, F., 2019. Decoupled weight decay regularization. In: *International Conference on Learning Representations*. <https://openreview.net/forum?id=Bkg6RiCqY7>.
- Mazal, P.P., Haehner, A., Hummel, T., 2016. Relation of the volume of the olfactory bulb to psychophysical measures of olfactory function. *Eur Arch Otorhinolaryngol* 273 (1), 1–7.
- McGraw, K.O., Wong, S.P., 1996. Forming inferences about some intraclass correlation coefficients. *Psychol Methods* 1 (1), 30.
- Merkel, D., 2014. Docker: lightweight linux containers for consistent development and deployment. *Linux journal* 2014 (239), 2.
- Milletari, F., Navab, N., Ahmadi, S.-A., 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: *2016 fourth international conference on 3D vision (3DV)*. IEEE, pp. 565–571.
- Mueller, A., Abolmaali, N., Hakimi, A., Gloeckler, T., Herting, B., Reichmann, H., Hummel, T., 2005. Olfactory bulb volumes in patients with idiopathic parkinson's disease: a pilot study. *J. Neural Transm.* 112 (10), 1363–1370.
- Mugler III, J.P., 2014. Optimized three-dimensional fast-spin-echo mri. *J. Magn. Reson. Imaging* 39 (4), 745–767.
- Noothout, J.M., Postma, E.M., Boesveldt, S., De Vos, B.D., Smeets, P.A., Išgum, I., 2021. Automatic segmentation of the olfactory bulbs in mri. In: *Medical Imaging 2021: Image Processing*, Vol. 11596. International Society for Optics and Photonics, p. 115961J.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A., 2017. Automatic differentiation in pytorch.
- Pérez-García, F., Sparks, R., Ourselin, S., 2020. Torchio: a python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning. *arXiv:2003.04696 [cs, eess, stat]*. ArXiv: 2003.04696. <http://arxiv.org/abs/2003.04696>.
- R Core Team, 2020. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org/>.
- Roberts, R.O., Christianson, T.J., Kremers, W.K., Mielke, M.M., Machulda, M.M., Vassilaki, M., Alhurani, R.E., Geda, Y.E., Knopman, D.S., Petersen, R.C., 2016. Association between olfactory dysfunction and amnesic mild cognitive impairment and alzheimer disease dementia. *JAMA Neurol* 73 (1), 93–101.
- Rombaux, P., Grandin, C., Duprez, T., 2009. How to measure olfactory bulb volume and olfactory sulcus depth? *Acta Otorhinolaryngol Belg* 8, 53.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. Springer, pp. 234–241.
- Roy, A.G., Conjeti, S., Navab, N., Wachinger, C., Initiative, A.D.N., et al., 2019. Bayesian quicknat: model uncertainty in deep whole-brain segmentation for structure-wise quality control. *Neuroimage* 195, 11–22.
- Roy, A.G., Conjeti, S., Navab, N., Wachinger, C., Initiative, A.D.N., et al., 2019. Quicknat: a fully convolutional network for quick and accurate segmentation of neuroanatomy. *Neuroimage* 186, 713–727.
- Roy, A.G., Conjeti, S., Sheet, D., Katouzian, A., Navab, N., Wachinger, C., 2017. Error corrective boosting for learning fully convolutional networks with limited data. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 231–239.
- Roy, A.G., Navab, N., Wachinger, C., 2018. Recalibrating fully convolutional networks with spatial and channel “squeeze and excitation” blocks. *IEEE Trans Med Imaging* 38 (2), 540–549.
- Seubert, J., Freiherr, J., Frasnelli, J., Hummel, T., Lundström, J.N., 2012. Orbitofrontal cortex and olfactory bulb volume predict distinct aspects of olfactory performance in healthy subjects. *Cerebral Cortex* 23 (10), 2448–2456.
- Shattuck, D.W., Leahy, R.M., 2002. BrainSuite: an automated cortical surface identification tool. *Med Image Anal* 6 (2), 129–142.
- Sinha, A., Dolz, J., 2020. Multi-scale self-guided attention for medical image segmentation. *IEEE J Biomed Health Inform.*
- Stöcker, T., 2016. Big data: the rhineland study. In: *Proceedings of the 24th Scientific Meeting of the International Society for Magnetic Resonance in Medicine (Singapore)*.

- Sudre, C.H., Cardoso, M.J., Ourselin, S., Initiative, A.D.N., et al., 2017. Longitudinal segmentation of age-related white matter hyperintensities. *Med Image Anal* 38, 50–64.
- Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Cardoso, M.J., 2017. Generalised Dice Overlap as a Deep Learning Loss Function for Highly Unbalanced Segmentations. In: *Deep learning in medical image analysis and multimodal learning for clinical decision support*. Springer, pp. 240–248.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9.
- Taha, A.A., Hanbury, A., 2015. Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. *BMC Med Imaging* 15 (1), 1–28.
- Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T.E., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., et al., 2012. The human connectome project: a data acquisition perspective. *Neuroimage* 62 (4), 2222–2231.
- Vassilaki, M., Christianson, T.J., Mielke, M.M., Geda, Y.E., Kremers, W.K., Machulda, M.M., Knopman, D.S., Petersen, R.C., Lowe, V.J., Jack Jr, C.R., Roberts, R.O., 2017. Neuroimaging biomarkers and impaired olfaction in cognitively normal individuals. *Ann. Neurol.* 81 (6), 871–882.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need. In: *NIPS*, pp. 6000–6010. <http://papers.nips.cc/paper/7181-attention-is-all-you-need>.
- Wang, J., You, H., Liu, J.-F., Ni, D.-F., Zhang, Z.-X., Guan, J., 2011. Association of olfactory bulb volume and olfactory sulcus depth with olfactory function in patients with parkinson disease. *American journal of neuroradiology* 32 (4), 677–681.
- Weiss, T., Soroka, T., Gorodisky, L., Shushan, S., Snitz, K., Weissgross, R., Furman-Haran, E., Dhollander, T., Sobel, N., 2020. Human olfaction without apparent olfactory bulbs. *Neuron* 105 (1), 35–45.
- Wilcoxon, F., 1992. Individual Comparisons by Ranking Methods. In: *Breakthroughs in statistics*. Springer, pp. 196–202.
- Yousem, D.M., Geckle, R.J., Bilker, W.B., Doty, R.L., 1998. Olfactory bulb and tract and temporal lobe volumes: normative data across decades a. *Ann. N. Y. Acad. Sci.* 855 (1), 546–555.
- Yousem, D.M., Geckle, R.J., Doty, R.L., Bilker, W.B., 1997. Reproducibility and reliability of volumetric measurements of olfactory eloquent structures. *Acad Radiol* 4 (4), 264–269.
- Zhang, H., Goodfellow, I., Metaxas, D., Odena, A., 2019. Self-attention generative adversarial networks. In: *International conference on machine learning*. PMLR, pp. 7354–7363.