



# FastSurfer-HypVINN: Automated sub-segmentation of the hypothalamus and adjacent structures on high-resolucional brain MRI

Santiago Estrada<sup>a,b</sup>, David Kügler<sup>a</sup>, Emad Bahrami<sup>a,c</sup>, Peng Xu<sup>b</sup>, Dilshad Mousa<sup>b</sup>, Monique M.B. Breteler<sup>b,d</sup>, N. Ahmad Aziz<sup>b,e</sup>, Martin Reuter<sup>a,f,g</sup>

<sup>a</sup>AI in Medical Imaging, German Center for Neurodegenerative Diseases (DZNE), Bonn, Germany

<sup>b</sup>Population Health Sciences, German Center for Neurodegenerative Diseases (DZNE), Bonn, Germany

<sup>c</sup>Computer Science Department, University of Bonn, Bonn, Germany

<sup>d</sup>Institute for Medical Biometry, Informatics and Epidemiology (IMBIE), Faculty of Medicine, University of Bonn, Bonn, Germany

<sup>e</sup>Department of Neurology, Faculty of Medicine, University of Bonn, Bonn, Germany

<sup>f</sup>A.A. Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Boston, MA, United States

<sup>g</sup>Department of Radiology, Harvard Medical School, Boston, MA, United States

Corresponding Author: Martin Reuter ([martin.reuter@dzne.de](mailto:martin.reuter@dzne.de))

## ABSTRACT

The hypothalamus plays a crucial role in the regulation of a broad range of physiological, behavioral, and cognitive functions. However, despite its importance, only a few small-scale neuroimaging studies have investigated its sub-structures, likely due to the lack of fully automated segmentation tools to address scalability and reproducibility issues of manual segmentation. While the only previous attempt to automatically sub-segment the hypothalamus with a neural network showed promise for 1.0 mm isotropic T1-weighted (T1w) magnetic resonance imaging (MRI), there is a need for an automated tool to sub-segment also high-resolucional (HiRes) MR scans, as they are becoming widely available, and include structural detail also from multi-modal MRI. We, therefore, introduce a novel, fast, and fully automated deep-learning method named *HypVINN* for sub-segmentation of the hypothalamus and adjacent structures on 0.8 mm isotropic T1w and T2w brain MR images that is robust to missing modalities. We extensively validate our model with respect to segmentation accuracy, generalizability, in-session test-retest reliability, and sensitivity to replicate hypothalamic volume effects (e.g., sex differences). The proposed method exhibits high segmentation performance both for standalone T1w images as well as for T1w/T2w image pairs. Even with the additional capability to accept flexible inputs, our model matches or exceeds the performance of state-of-the-art methods with fixed inputs. We, further, demonstrate the generalizability of our method in experiments with 1.0 mm MR scans from both the Rhineland Study and the UK Biobank—an independent dataset never encountered during training with different acquisition parameters and demographics. Finally, *HypVINN* can perform the segmentation in less than a minute (graphical processing unit [GPU]) and will be available in the open source *FastSurfer* neuroimaging software suite, offering a validated, efficient, and scalable solution for evaluating imaging-derived phenotypes of the hypothalamus.

**Keywords:** hypothalamic sub-segmentation, deep learning, hetero-modal, high-resolution, structural MRI

## 1. INTRODUCTION

### 1.1. Motivation

The hypothalamus consists of a group of interconnected neuronal nuclei located at the base of the brain (Saper &

Lowell, 2014). It is the body's principal homeostatic center and plays a crucial role in the regulation of a broad range of physiological, behavioral, and cognitive functions, both through direct control of endocrine and autonomic nervous system outflow, as well as through

Received: 20 June 2023 Revision: 26 September 2023 Accepted: 1 November 2023 Available Online: 6 November 2023



© 2023 Massachusetts Institute of Technology.  
Published under a Creative Commons Attribution 4.0  
International (CC BY 4.0) license.

Imaging Neuroscience, Volume 1, 2023  
[https://doi.org/10.1162/imag\\_a\\_00034](https://doi.org/10.1162/imag_a_00034)

extensive projections to cortical and limbic regions (Saper & Lowell, 2014). Neuropathological studies have demonstrated extensive involvement of the hypothalamus in a range of neurodegenerative diseases, including Alzheimer's disease (Liguori et al., 2014; Roh et al., 2014), Parkinson's disease (Fronczek et al., 2007), Huntington's disease (van Wamelen & Aziz, 2021), frontotemporal dementia, and amyotrophic lateral sclerosis (Ahmed et al., 2021; Bocchetta et al., 2015). However, the association between hypothalamic integrity and physiological, behavioral, and cognitive outcomes has not been studied in large clinical or population-based studies for lack of a reliable high-throughput automatic imaging procedure.

The majority of studies on hypothalamic imaging-derived phenotypes use manual annotations of magnetic resonance imaging (MRI) scans as the gold standard. Manual segmentation of the hypothalamus and its substructures is commonly done on T1-weighted images (Makris et al., 2013; Schindler et al., 2013). Nonetheless, the use of multi-modal structural information during the manual annotation process has also been proposed to increase especially the visibility of the lateral hypothalamus boundaries (Baroncini et al., 2012; Bocchetta et al., 2015). These multi-modal protocols recommend segmenting the hypothalamus using simultaneous visualization of registered T1-weighted (T1w) and T2-weighted (T2w) MR images. Manual delineation of the hypothalamus, however, is a very time-consuming process that relies highly on the user's expertise due to the small size and low boundary MR contrast in the hypothalamus region, regardless of the available MRI modalities.

Automated methods have been proposed to segment the whole hypothalamus (Greve et al., 2021; Orbes-Arteaga et al., 2015; Rodrigues et al., 2020, 2022; Thomas et al., 2019) and its sub-regions (Billot, Bocchetta, et al., 2020) quickly and reliably. However, even though automated tools are available, they only focus on segmenting 1.0 mm isotropic T1w scans, ignoring the detailed structural information available in sub-millimeter resolution datasets. High-resolutional (HiRes) MR scans are becoming more common across studies (even in clinical settings) due to rapid advancements in MR technology (e.g., accelerated acquisition schemes) and are increasingly employed as the new standard for large studies (e.g., the Rhineland Study (Breteler et al., 2014; Stöcker, 2016), Human Connectome Project (HCP) datasets (Bookheimer et al., 2019; Harms et al., 2018; Van Essen et al., 2012), Autism Brain Imaging Data Exchange II (ABIDE-II) (Di Martino et al., 2017), and TRACK-PD (Wolters et al., 2020)). Thus, the

need for neuroimaging tools that can handle sub-millimeter resolutions (e.g., 0.8 mm isotropic) has increased.

Moreover, current automated hypothalamic segmentation methods have neglected the inclusion of multi-modal structural information. One reason for this is that simultaneous access to T1w and T2w images is not always possible due to constraints in scanning time or poor image quality in one of the modalities due to reduced image resolution or acquisition artifacts. Therefore, the introduction of an accurate automated method for segmenting hypothalamic structures on high-resolutional T1w and T2w MRI scans, which is also robust to handle missing modalities, is of significant interest to clinicians and researchers.

## 1.2. Related work

Automated hypothalamic segmentation methods utilizing multi-atlas-based techniques (Orbes-Arteaga et al., 2015; Thomas et al., 2019) were initially proposed. However, these methods are slow and demand considerable computational resources. Newer techniques such as fully convolutional neural networks (F-CNNs) can tremendously speed up computation time by utilizing graphical processing units (GPUs) and have become the preferred method for solving supervised semantic segmentation problems in the medical computer vision community (Estrada et al., 2020, 2021; Faber et al., 2022; Henschel et al., 2020; Kamnitsas et al., 2017; Milletari et al., 2016; Ronneberger et al., 2015; Roy et al., 2019).

Hypothalamus segmentation using F-CNNs has mainly focused on identifying the hypothalamus as one whole structure in the brain (Greve et al., 2021; Rodrigues et al., 2020, 2022). Recently, Billot, Bocchetta, et al. (2020) proposed a method to segment five sub-regions of the hypothalamus using an encoder-decoder 3D F-CNN with extensive data augmentation. They followed the hypothalamic parcellation protocol introduced by Makris et al. (2013) on standard 1.0 mm isotropic resolution T1w images. Their proposed method illustrates the capabilities of F-CNNs to segment hypothalamic compartments with promising results on datasets acquired at 1.0 mm isotropic resolution (Billot, Bocchetta, et al., 2020; Shapiro et al., 2022). However, F-CNNs are known to have issues generalizing to resolutions that differ from the training one (Estrada et al., 2021; Henschel et al., 2022; Iglesias et al., 2021), rendering HiRes images out-of-distribution and unsuitable for methods designed for lower resolutions. A common approach for this problem is to down-sample the input image to the desired lower resolution in a pre-processing step (Billot, Bocchetta, et al., 2020; Greve

et al., 2021; Henschel et al., 2020). This process, however, reduces image details and information, forfeiting the investment already made when acquiring the higher resolution in the first place. Furthermore, HiRes information could help address inter-class inconsistencies between voxels at a local and global level and alleviate the partial volume effect problem (Glasser et al., 2013).

HiRes segmentation of brain structures has mostly been tackled by training with manual annotations created at the desired resolution (Beliveau et al., 2021; Estrada et al., 2021; Kamnitsas et al., 2017; Rushmore et al., 2022) or training models using 1.0 mm data with scale augmentations—an established deep-learning technique to improve the generalizability of a model. Recently, models capable of segmenting scans at different resolutions have been introduced. Billot, Colin, et al. (2023) and Billot, Greve, et al. (2023) proposed *SynthSeg*, a technique for generating segmentations at a fixed resolution (1.0 mm), regardless of the resolution of the input scan, which are interpolated to the fixed resolution as a pre-processing step. During training, *SynthSeg* relies on a generative model that produces “unrealistic synthetic images” (Billot, Greve, et al., 2023). These synthetic images are created from ground truth label maps at the pre-defined fixed resolution. This approach simulates domain variability by incorporating multiple random parameters for the generator, such as spatial, intensity, contrast, and resolution variability. While providing input flexibility, the model’s output resolution, however, remains confined to the fixed resolution.

Before *SynthSeg*, we introduced the Voxel-Size Independent Neural Network (VINN) for resolution-independent segmentation tasks (Henschel et al., 2022). The VINN approach enables training and inference using images at multiple resolutions within a single network. In brief, instead of interpolating input images, VINN integrates the resolution change into the network, replacing a regular scale transition with an interpolation layer that maps the latent space at native input resolution to a pre-defined internal resolution at lower layers of the network and vice versa. As a result, rich HiRes information is retained without image or label interpolation, and segmentations are provided at the desired native input resolution.

Finally, as has already been shown in manual segmentation of hypothalamic structures, exclusively utilizing T1w images as input forfeits the significant potential presented by the inclusion of multi-modal information (T1w and T2w) (Baroncini et al., 2012; Bocchetta et al., 2015). Common multi-modal F-CNN architectures, however, require all input modalities to always be present. The

absence of any modality introduces a computational bias that the network is not trained to handle. To overcome missing modalities, proposed solutions include training a specific network for each of the input combinations or providing the segmentation model with a synthesized version of unavailable modalities (Hofmann et al., 2008; Van Tulder & de Bruijne, 2015). Alternatively, training networks with synthetic image contrast has also been suggested (Billot, Greve, et al., 2020, 2023). Even though these techniques have shown promising results, a more suitable model should be capable of extracting the most salient information for solving the given task from the available modalities without the need for artificial images or multiple modality-specific networks. With this in mind, shared latent space models were introduced on the challenging task of multi-modal brain tumor segmentation (Dorent et al., 2019; Havaei et al., 2016; Varsavsky et al., 2018). This approach first translates modalities into independent latent spaces; afterwards, the modalities’ embedded information is merged inside the network into a shared latent representation. The shared latent space is then forwarded to the remaining network to solve the desired task. At inference time, the shared representation is computed from the available modalities, thus being robust to all input-modality combinations (i.e., hetero-modal) included in training.

To address the missing modalities challenge in an HiRes scenario, we suitably include the shared latent space concept into our voxel-size independent network (VINN). *Hetero-modal VINN* (HM-VINN) introduces a fusion module that linearly combines the modalities inside the network. After passing the available scans through a separate modality-specific convolutional block, the network weighs and merges the feature maps based on the best available information using a learnable weighted sum. As the output of the fusion module is normalized, missing one modality can be tackled by assigning zero to its respective weight.

### 1.3. Contribution

To our knowledge, we are the first to tackle automated hetero-modal sub-segmentation of the hypothalamus and adjacent structures on high-resolution brain MRI. The contributions of this work are the following: Firstly, we introduce a new hypothalamic labeling protocol adapted to the higher spatial resolution offered by 3 T 0.8 mm isotropic MR images. The proposed protocol presents a more fine-grained parcellation of the hypothalamus and includes usually ignored brain structures, such as hypophysis,

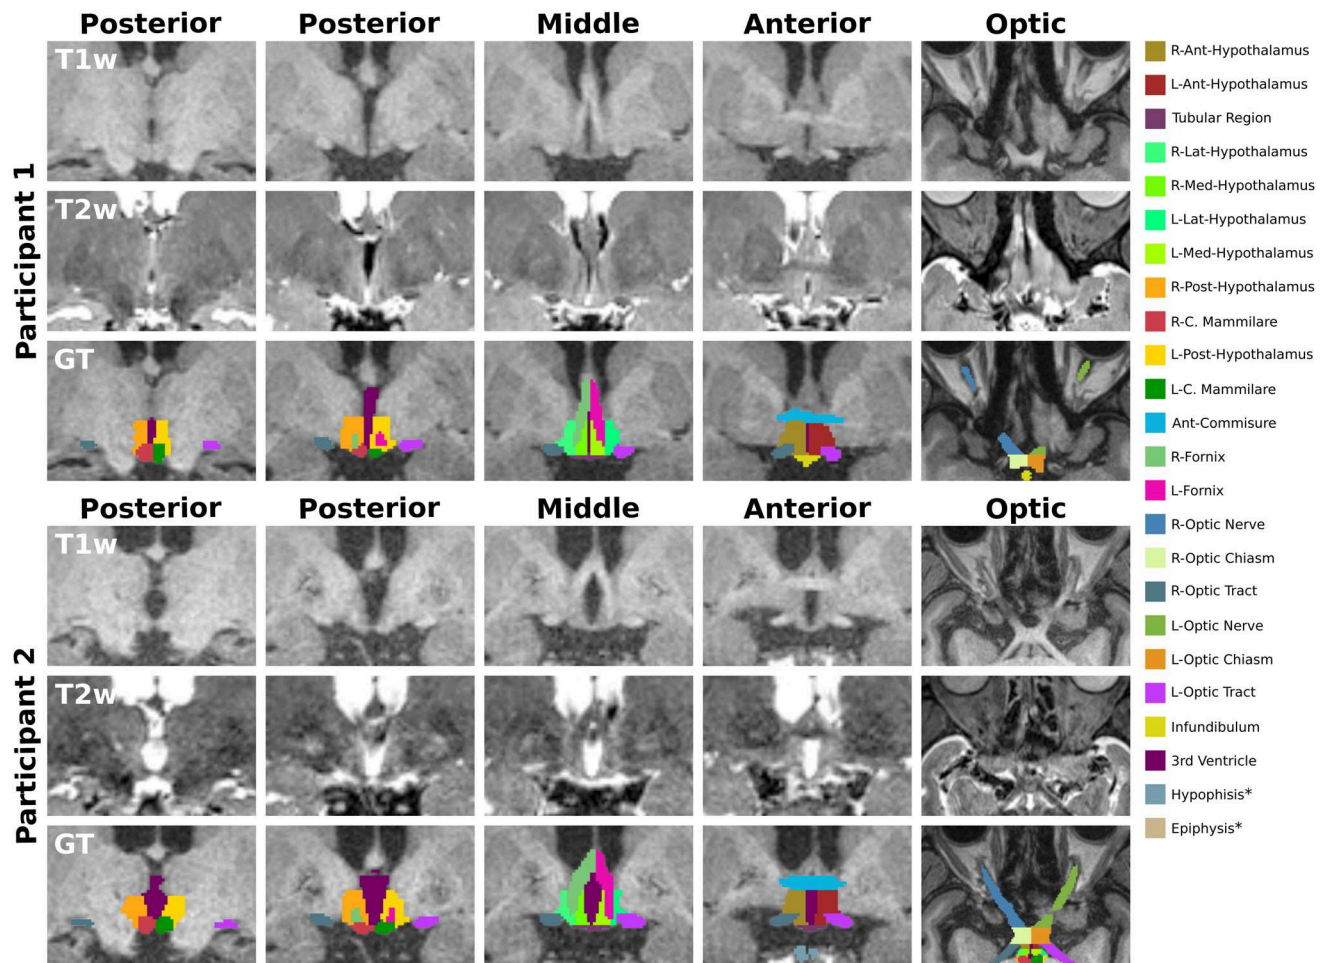
epiphysis, the optic nerve, optic chiasm, and optic tract, as illustrated in Figure 1. Secondly, we present *HypVINN*, a novel automated hypothalamic parcellation tool with a novel *hetero-modal VINN (HM-VINN)* architecture at its core, providing a solution to the multi-resolution and the missing modality challenge in a single model. We extensively show that the model's input flexibility does not compromise performance compared to state-of-the-art methods with fixed inputs in terms of segmentation accuracy, test-retest reliability, and generalizability. Moreover, our method replicates hypothalamic volume effects (e.g., age and sex) on subsets of the 0.8 mm (HiRes) Rhineland Study ( $n = 463$ ) and the 1.0 mm UK Biobank ( $n = 535$ ) (Alfaro-Almagro et al., 2018; Miller et al., 2016). Last but not least, and to the benefit of the research community, we will integrate the *HypVINN* tool into the user-friendly, open

source *FastSurfer* framework (Henschel et al., 2020) available at: <https://github.com/Deep-MI/FastSurfer> (code will be released upon acceptance).

## 2. METHODS

### 2.1. Datasets

We used MR images from two population studies, namely the Rhineland Study (RS) (Breteler et al., 2014; Stöcker, 2016) and the UK Biobank (UKB) (Alfaro-Almagro et al., 2018; Miller et al., 2016), with resolutions of 0.8 mm (HiRes) and 1.0 mm, respectively. Participants from both studies gave written informed consent in accordance with the ethical guidelines of the individual studies. Furthermore, ethics approval and regulations can be accessed on their



**Fig. 1.** T1-weighted (T1w) and T2-weighted (T2w) images and ground truth (GT) from two participants. The proposed manual segmentation scheme is composed of twenty-four structures divided into three major regions: 1) hypothalamic (anterior, middle, and posterior), 2) optic, and 3) others. The color lookup table\* for all structures is presented on the left, and a detailed overview of the three regions is presented in Table 1. \*Structures are not visible in the presented snapshots.



**Table 1.** Summary of the hypothalamic sub-regions and adjacent structures included in the proposed labeling scheme with its corresponding name, anatomical designation, and region.

Hypothalamic sub-regions			Adjacent structures		
Label name	Anatomical designation	Region group	Label name	Anatomical designation	Region group
L-Ant-Hypothalamus	Anterior Hypothalamus (lh), Supraoptic Nucleus (lh)	Anterior	3rd-Ventricle	3rd-Ventricle, Superior-Border	Others
R-Ant-Hypothalamus	Anterior Hypothalamus (rh), Supraoptic Nucleus (rh)		L-Fornix R-Fornix	Fornix (lh) Fornix (rh)	
L-Med-Hypothalamus	Medial Hypothalamus* (lh)		Epiphysis Hypophysis	Epiphysis Hypophysis, Neurohypophysis	
R-Med-Hypothalamus	Medial Hypothalamus* (rh)	Middle	Infundibulum	Infundibulum	Optic
L-Lat-Hypothalamus	Lateral-Hypothalamus (lh)		Ant-Commisure	Anterior Commisure	
R-Lat-Hypothalamus	Lateral-Hypothalamus (rh)		L-N-Opticus R-N-Opticus	Optic Nerve (lh) Optic Nerve (rh)	
Tuberal-region	Median-eminence, Tuberomammillary Region, and Arcuate-nucleus		L-Chiasma-Opticus R-Chiasma-Opticus	Optic Chiasm (lh) Optic Chiasm (rh)	
L-Post-Hypothalamus	Posterior Hypothalamus (lh)	Posterior	L-Optic-tract R-Optic-tract	Optic Tract (lh) Optic Tract (rh)	
R-Post-Hypothalamus	Posterior Hypothalamus (rh)				
L-C-Mammillare	Corpus Mammillare (lh)				
R-C-Mammillare	Corpus Mammillare (rh)				

\*Including the Paraventricular Nucleus (PVN), the Ventromedial Nucleus (VMN), and the Dorsomedial Nucleus (DMN).

respective webpages. For this work, we compiled four distinct datasets from the population studies: a manually annotated dataset (from RS), a generalizability dataset (from RS and UKB), a test-retest dataset (from RS), and a case-study dataset (from RS and UKB). The manually annotated dataset (referred to as "in-house dataset") was initially split into two non-overlapping sets, one for training and validation, and the other for testing. The remaining datasets were exclusively used for evaluations to assess different aspects of our hetero-modal method.

The Rhineland Study is an ongoing population-based cohort study located in Bonn, Germany, which enrolls participants aged 30 years and above ([www.rheinland-studie.de](http://www.rheinland-studie.de)). MR scans were collected at two different sites using identical 3 T Siemens MAGNETOM Prisma MRI scanners equipped with 64-channel head-neck coils. The core MRI acquisition protocol for every participant in the Rhineland Study includes the following MR contrast: T1w, T2w, FLAIR, diffusion-weighted, susceptibility-weighted, resting-state functional, and abdominal Dixon MRI with a total net scan time of around 45 minutes. Furthermore, an optional extra acquisition time (maximum 10 minutes) is available for a free protocol.

This paper utilized the 0.8 mm isotropic T1w and T2w MR scans. The T1 protocol consists of a multi-echo magnetization prepared rapid gradient echo (MPRAGE) sequence ([van der Kouwe et al., 2008](#)) with 2D acceleration ([Brenner et al., 2014](#)), while the T2 protocol uses a 3D Turbo-Spin-Echo (TSE) sequence with variable flip angles ([Busse et al., 2008](#)). Both sequences also utilize elliptical sampling ([Mugler III, 2014](#)) and parallel imaging (PI) ([Griswold et al., 2002](#)) to expedite the imaging process. For this work, all protocol versions from the Rhineland Study were considered, and sequence parameters are presented in [Appendix Table A1](#).

We compiled the Rhineland Study datasets by first randomly selecting a subset ( $n = 534$ ) of participants with available T1w and T2w scans from sex and age strata to ensure a balanced population distribution. The sample presents a mean age of 54.9 years (range 30 to 95), and 59.4% were women. We then further assigned participants to the in-house dataset and all its subsequent splits adhering to the age and sex-stratification scheme. All T2w scans were registered to their corresponding T1w scan using *FreeSurfer's* *mri\_robust\_register* tool ([Reuter et al., 2010](#)).

MRI scans of the *in-house training and testing dataset* ( $n = 50$ ) were manually annotated by an experienced rater and split into training/validation ( $n = 44$ ) and testing ( $n = 6$ ) sets. Training data were further split into four groups for cross-validation. Finally, the testing data were manually annotated for a second time by our main rater to evaluate intra-rater variability. The rater was blind to the scans' identification to avoid bias and overestimating performance.

For evaluating within-session *test-retest* reliability, we utilized the RS subset ( $n = 21$ ) with two in-session T1w scans. The additional scan for this participant was acquired during the time slot allocated for a free protocol inside the Rhineland study's MRI acquisition protocol. Due to the time constraint of the free protocol, a second T2w scan could not be acquired. Before starting the free protocol, participants were asked to move their head inside the head-neck coil. It is important to note that T1w scans were not acquired back-to-back, but with a time gap of almost 30 minutes.

The MRI scans of the remaining participants ( $n = 463$ ) were compiled into the *RS case-study dataset* to evaluate the sensitivity to known hypothalamic volume effects (e.g., age and sex). For a detailed description of the population characteristics of all the aforementioned RS subsets, see [Appendix Tables A2 and A3](#).

We used data from the UK Biobank study to test the *generalizability* of our method to isotropic 1.0 mm scans from an unseen cohort with different acquisition parameters. An initial subset ( $n = 544$ ) of random participants was selected from sex and age strata to ensure a balanced population distribution. The chosen sample presents a mean age of 58.7 years (range 45 to 82), consisting of 52.6% women. Subsequently, the scans of nine random participants were manually labeled by our expert rater to evaluate segmentation accuracy at 1.0 mm (*generalizability dataset*). The remaining UKB participants ( $n = 535$ , *UKB case-study dataset*) were also used in the hypothalamic volumes effects sensitivity analysis. A summary of the population characteristic of the UKB subsets is presented in [Appendix Table A4](#).

## 2.2. Manual reference standard

An experienced rater manually annotated the sub-regions of the hypothalamus and adjacent structures on registered T1w and T2w images, except for the UK Biobank cases where only T1w scans were available. The annotation was performed using *Freeview*, a visualization tool of *FreeSurfer* ([Fischl, 2012](#); [Fischl et al., 2002](#)), which allowed simultaneous viewing of the available modalities.

Summarizing the labeling process, the borders of the unilateral hypothalamus were defined as follows ([Makris et al., 2013](#)): a) anteriorly: coronal plane passing through the most rostral tip of the anterior commissure and containing the optic chiasm, b) posteriorly: coronal plane through the most caudal tip of the mammillary bodies, c) superiorly: third ventricle with the diencephalic fissure, d) inferiorly: junction to the optic chiasm rostrally and the hemispheric margin more caudally, e) medially: wall of the third ventricle and the interhemispheric fissure, and f) laterally: rostrally at the medial border of the optic tract and more caudally at the internal capsule, globus pallidus, and cerebral penduncle. A detailed definition of the segmentation procedure for all different substructures is provided in [Appendix C](#). Adjacent small hypothalamic nuclei were grouped into subunits according to [Table 1](#). An example of the manual segmentation scheme is illustrated in [Figure 1](#), and an overview of all twenty-four segmented structures is presented in [Table 1](#).

## 2.3. Hypothalamic hetero-modal segmentation tool—HypVINN

### 2.3.1. Hetero-modal segmentation network—HM-VINN

To accurately segment the hypothalamic sub-regions and adjacent structures, we employ *VINN* ([Henschel et al., 2022](#)) as the foundation for our network design. *VINN* is a resolution-independent extension of the successful multi-network approach *FastSurferCNN* ([Estrada et al., 2021](#); [Faber et al., 2022](#); [Henschel et al., 2020](#)). Both methods are 2.5D approaches, that is, they aggregate predictions of three 2D F-CNNs (one per anatomical view) with multi-slice input ([Henschel et al., 2020](#)). The F-CNNs follow a UNet-type layout with an encoder and decoder arm of five competitive-dense blocks (CDB) separated by an additional bottleneck CDB (see [Fig. 2](#)). In *FastSurferCNN*, all scale transitions between the CDBs are implemented via fixed-scale down- or up-sampling operations (i.e., (un) pooling). *VINN*, on the other hand, replaces the first and last scale transition with a flexible network-integrated resolution-normalization. Here, the native image resolution is explicitly integrated into the network and utilized to interpolate the feature maps to a common pre-defined network base resolution (1.0 mm). In turn, network capacity in the inner layers is available for the segmentation task while retaining voxel size-dependent information outside of it. Lastly, the view-aggregation step ensembles the resulting probabilities maps through a weighted average (axial = 0.4, coronal = 0.4, and sagittal = 0.2). The weights of the sagittal predictions are reduced compared to

the other predictions, as structures with left and right hemispheres labels are unified into one due to missing lateralization information in the sagittal view (Henschel et al., 2020). For the current segmentation task, we also unify lateralized structure labels into one for the sagittal view, consequently reducing the number of classes in the sagittal F-CNN from 24 to 15. Therefore, the *VINN* view-aggregation weighting scheme is also suitable for our application.

In this work, we extend *VINN* into a hetero-modal segmentation scenario (referred to as *HM-VINN*) by embedding the input modalities into a shared latent space (Dorent et al., 2019; Havaei et al., 2016; Varsavsky et al., 2018). Following this direction, we modify the standard F-CNNs from *VINN* to initially process T1w and T2w images independently of each other by replacing the first encoder CDB with modality-specific CDBs (Fig. 2, e.g., T1-CDB\* and T2-CDB\*). After the independent stage, feature maps are merged inside the network by a fusion module and fed into the following convolutional pipeline.

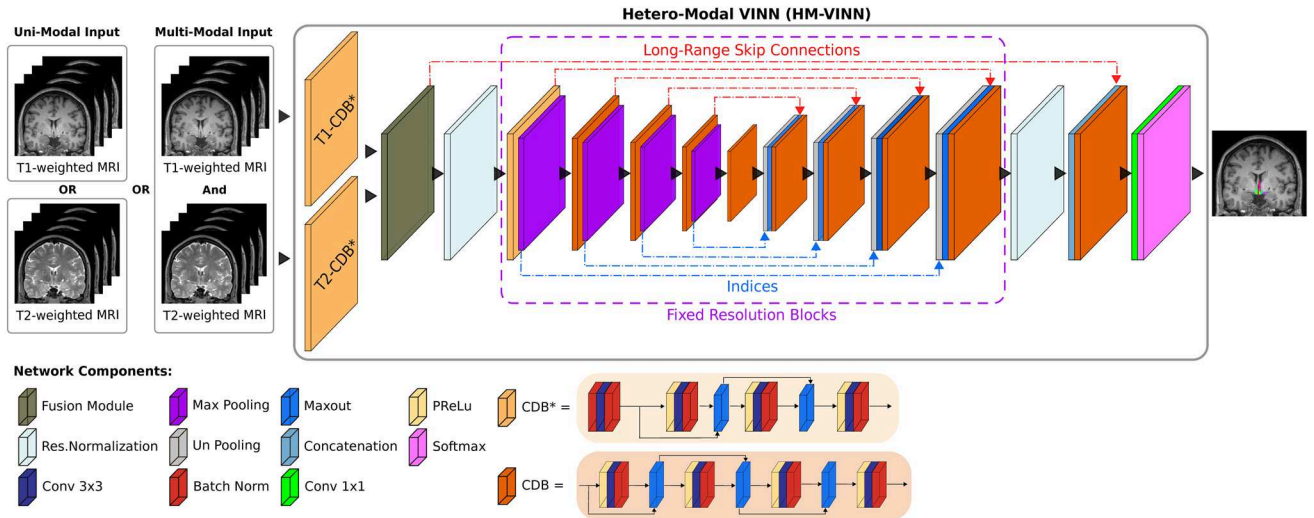
The implemented fusion module weights and merges the feature maps from the T1 and T2 branches based on the best available information using a learnable weighted sum. Let us denote the output feature map from the T1-CDB\* as  $F_{T1} \in \mathbb{R}^{C \times H \times W}$  and the T2-CDB\* output as  $F_{T2} \in \mathbb{R}^{C \times H \times W}$ , where  $C, H, W$  represent the channel, height, and width

dimensions, respectively. Then, the output of fusion module  $F_{\text{fused}}$  is

$$F_{\text{fused}} = \frac{|W_{T1}|}{|W_{T1}| + |W_{T2}|} \times F_{T1} + \frac{|W_{T2}|}{|W_{T1}| + |W_{T2}|} \times F_{T2}, \quad (1)$$

where  $W_{T1}$  and  $W_{T2}$  are global learnable scalar parameters initialized both at 0.5. The introduction of  $W_{T1}$  and  $W_{T2}$  allows the network to gradually learn the importance of each modality. If a modality is more informative, its feature maps will have a higher weight. Additionally, as the output of the fusion module is normalized, missing one modality can be tackled by assigning zero to its respective weight. Thus, the fusion features are identical to the encoder block output of the existing modality.

In detail, all three F-CNNs followed the abovementioned layout (see Fig. 2). Within F-CNNs, the CDB layout is kept mostly the same as the one from *VINN*, where the CDB consists of four layers of parametric rectified linear unit (PReLU), convolution (Conv - kernel size of  $3 \times 3$ ), and batch normalization (BN) except for the first two encoders blocks. In the first two encoder blocks from *VINN*, the first PReLU is replaced with a BN to normalize the inputs (see Fig. 2, CBD\*). The modified CBD construction is also utilized for modality-specific CDBs as they are our initial first encoder CDB. To keep the



**Fig. 2.** *Hetero-Modal VINN (HM-VINN)* architecture in *HypVINN*. Input modalities are first independently processed by modality-specific competitive dense blocks (T1-CDB\* and T2-CDB\*). Afterward, modality-specific feature maps are merged inside the network by our proposed fusion module (dark green) to create a shared latent space. During inference time, the shared latent space can be computed over the available modalities and fed into the remaining network. Furthermore, *HM-VINN* incorporates flexible transitions in the first and last scale transition by utilizing the network-integrated resolution-normalization (light blue). Each CDB is composed of four sequences of parametric rectified linear unit (PReLU), convolution (Conv), and batch normalization (BN). In the modality-specific CDBs and second encoder block (CBD\*), the first PReLU is replaced with a BN to normalize the inputs.

comparison fair in light of an effective parameter count of approximately 4.5 M parameters (three dedicated, modality-specific models with approx. 1.5 M parameters each), we increase the number of channels (features) of all layers from 64 to 80 inside CDBs, and from 32 to 64 in the first and last CDB blocks (i.e., the first scale level). This change raises the parameter count to approximately 2.6 M, which is still significantly less than three ( $\approx 4.5$  M parameters) or even two ( $\approx 3.0$  M) dedicated, modality-specific networks.

### 2.3.2. Hetero-modal training procedure

Introducing additional variations by data augmentation during training helps neural networks to be more robust. Here, we make *HM-VINN* robust to missing modalities by sometimes randomly dropping either the T1w or T2w image for a given training example with a uniform distribution between all input combinations (modality dropout). The modality weights in the fusion module are adjusted as follows: i) When the two modalities are available, the network automatically assigns the weights (see Eq. 1). ii) If a modality is dropped, its corresponding fusion weight is set to zero as described in the previous section. By starting this modality dropout procedure only after 10 epochs, the proposed training procedure first establishes general segmentation capabilities (with all modalities available) before pivoting to more difficult scenarios with different combinations and missing modalities.

### 2.3.3. Model learning

All F-CNN are implemented in PyTorch (Paszke et al., 2017) using a docker container (Merkel, 2014). Independent models for axial, coronal, and sagittal views are trained for 100 epochs with a batch size of 16 using two NVIDIA Tesla V100 GPU with 32 GB RAM. We use the AdamW (Kingma & Ba, 2015; Loshchilov & Hutter, 2019) optimizer with a weight decay of  $10^{-4}$  and an initial learning rate of 0.05, which is decreased to 0.005 after 70 epochs. The networks are trained by optimizing a combined loss function of a median frequency-weighted cross-entropy loss and Dice loss (Roy et al., 2019). This loss function encourages correct segmentation along anatomical boundaries and counters class imbalances by increasing the weights of less frequent classes.

To increase the generalizability of our model, we apply several spatial and intensity data augmentations during training. Spatial augmentations on the inputs images are limited to random affine transformations such as transla-

tion (range: from  $-15$  mm to  $15$  mm), rotation (range: from  $-10^\circ$  to  $10^\circ$ ), and uniform scaling (factor: from 0.85 to 1.15) (Pérez-García et al., 2020). Furthermore, we include internal scale augmentations of the feature maps as introduced by *FastSurferVinn* to improve the segmentation performance (Henschel et al., 2022).

Intensity augmentations are carried out to address two challenges: 1) intensity inhomogeneities due to scan parameters (Pérez-García et al., 2020) and 2) artefacts introduced by defacing algorithms in regions of interest (e.g., optic region). The first problem is tackled by applying a random bias field (Sudre et al., 2017; Van Leemput et al., 1999) transformation on the input images (coefficients range: from  $-0.5$  to  $0.5$ ). For the second issue, we improve the network's robustness to handle defaced scans by including scans with or without face features as part of the training set. For creating the modified scans, three common open-source algorithms are used (PyDeface (Gulban et al., 2019), MiDeFace from *FreeSurfer* (Fischl, 2012), and HCP face masking (Milchenko & Marcus, 2013)). In contrast to all above-mentioned transformations, defacing is performed statically before training ("offline") due to the high computation time to deface a scan (more than 1 minute per method).

## 2.4. Evaluation metrics

We compute three standard segmentation metrics (dice similarity coefficient, volume similarity, and Hausdorff distance) to assess the similarity between the predicted label maps and manual annotations (Taha & Hanbury, 2015). We first evaluate the dice similarity coefficient (Dice) (Dice, 1945; Sorensen, 1948) as it provides spatial overlap consensus. Let  $M$  (manual annotations) and  $P$  (prediction) denote binary label maps, then Dice is defined as:

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

where  $|M \cap P|$  represent the number of common elements (intersection) and  $|M|$  and  $|P|$  the number of elements in each label map; therefore, Dice values range from 0 to 1, and a higher Dice represents a better segmentation agreement. Afterwards, we compute volume similarity (VS) as volume measurements are usually the desired image-derived phenotype for downstream statistical analysis. VS is defined as:

$$VS = 1 - \frac{||M| - |P||}{|M| + |P|} \quad (3)$$



VS has the same range as Dice; however, it can have its maximum value even when the spatial overlap is zero, as this metric does not consider spatial localization information. Additionally, a spatial distance-based metric is used to evaluate the quality of segmentation boundary delineation (contour). Here, we use the 95% Hausdorff distance (HD95), a Hausdorff distance (HD) as it is less sensitive to outliers (Huttenlocher et al., 1993). HD95 is considered as the 95th percentile of the ordered distance measures, and it is defined as:

$$d_{95}(M, P) = 95^{th}_{m \in M} \left( \min_{p \in P} d(m, p) \right)$$

$$d_{HD95}(M, P) = \max(d_{95}(M, P), d_{95}(P, M)) \quad (4)$$

where  $d$  is the Euclidean distance. In contrast to the Dice and VS, HD95 is a dissimilarity metric so a smaller value indicates a better boundary delineation with a value of zero being the minimum (perfect match).

Finally, statistical significant differences in segmentation performance are confirmed throughout this work by a non-parametric paired two-sided Wilcoxon signed-rank test (Wilcoxon, 1992) after correcting for multiple testing using Bonferroni correction (referred to as corrected  $p$ ).

For accessing the test-retest reliability of predicted volume measurements between two repeated scans of the same participant, we use the intra-class correlation (ICC). The ICC is a commonly used metric to assess the degree of agreement and correlation between measurements. The ICC values range from 0 to 1, with higher values representing better reliability. Here, we compute a two-way fixed, absolute agreement and single measures with a 95% confidence interval (ICC(A,1)) (McGraw & Wong, 1996).

### 3. EXPERIMENTS AND RESULTS

This section is divided into four parts with the aim to thoroughly validate our hetero-modal method for hypothalamic sub-regions and adjacent structures segmentation (referred to as *HypVINN*). The *HypVINN* model is composed of the *HM-VINN* architecture and learning strategies introduced in Sections 2.3.2 and 2.3.3. i) Initially, we evaluate the segmentation **accuracy** of *HypVINN*'s predictions against manual annotations. For this purpose, we benchmark the network based on the performance in the unseen test-set against multi- and uni-modal models, including the only other contemporary method for hypothalamus parcellation (Section 3.1.1), and manual rater variability (Section 3.1.2). ii) We assess the **generalizabil-**

**ity** of our method to a different image resolution—1.0 mm isotropic MRI scans (Section 3.2). iii) We test the **reliability** of the predicted volumes in a within-session test-retest scenario (Section 3.3). iv) Finally, we measure the **sensitivity** of the proposed pipeline to replicate known hypothalamic volume effects with respect to age and sex. In order to ensure that all experiments are carried out under the same testing conditions, all inference analyses are evaluated in a Docker container with a 12 GB NVIDIA Titan V GPU. Model inference can also run on the CPU at reduced speeds.

#### 3.1. Accuracy

In this section, we benchmark and evaluate the accuracy of the hetero-modal *HypVINN*. All implemented networks are trained using the scheme mentioned in Section 2.3.3.

To show a proof-of-concept for our proposed *HypVINN* in segmenting hypothalamic sub-regions and adjacent structures with missing input modalities, we benchmark our method against segmentation scenarios where all modalities are always available (i.e., uni-modal and multi-modal models). For this purpose, we implement the classic *VINN* with three different inputs: i) only T1w (*T1-VINN*), ii) only T2w (*T2-VINN*), and iii) T1w & T2w (*multi-modal (MM)-VINN*). For the multi-modal model, the input passed to the network consists of a multi-channel image created by stacking T1w and T2w image slices on top of each other. Additionally, we compare our *HypVINN* against the method proposed by Billot, Bocchetta, et al. (2020)—a *3D-UNet* with extensive data augmentation for hypothalamic sub-segmentation on T1w images. Direct comparison of our predicted outcomes with the results from the already trained model from Billot et al. is not possible as our annotation protocol segments more structures and uses a different hypothalamic parcellation. Therefore, we utilize the implementation provided by the authors to retrain their T1w model from scratch with our manual annotations. It is important to notice that we do not fine-tune the implementation from Billot et al., and any optimization of their tool is outside this paper's scope. Furthermore, all comparative *VINN* baselines follow the same 2.5D scheme as mentioned in Section 2.3.1, and inference in *HypVINN* is done per input combination. The difference between results in the following two sections is in the data used for training: For Section 3.1.1 and Table 2, all networks are trained in a 4-fold cross-validation scheme to also generate validation performance on the holdout validation split (see Appendix B for ablation results). For all other results, we used the full training set

**Table 2.** Mean (and standard deviation) segmentation performance of the cross-validated F-CNN models on the unseen test-set.

Experimental setup	Dice $\uparrow$		VS $\uparrow$		HD95 (mm) $\downarrow$	
Model	Mean (SD)	Signif.	Mean (SD)	Signif.	Mean (SD)	Signif.
Only T1w input						
a: T1-VINN (Henschel et al., 2022)	0.7937 (0.0926)	c,d,e	0.9030 (0.0785)	c,e	1.1262 (0.5443)	c,d,e
b: HypVINN (Ours)	0.7905 (0.0968)	c,d,e	0.9053 (0.0757)	c,d,e	1.1312 (0.5683)	c,d,e
c: 3D-UNET (Billot, Bocchetta, et al., 2020)	0.7481 (0.1516)	d,e	0.8753 (0.1325)	e	1.4088 (2.235)	e
Only T2w input						
d: T2-VINN (Henschel et al., 2022)	0.7457 (0.1059)	e	0.8967 (0.0877)	c,e	1.2275 (0.5525)	e
e: HypVINN (Ours)	0.7224 (0.1120)		0.8683 (0.1074)		1.4315 (1.7678)	
Multi-modal (MM) input (T1w & T2w)						
f: MM-VINN (Henschel et al., 2022)	0.7918 (0.0924)	c,d,e	0.9033 (0.0774)	c,e	1.1350 (0.5819)	c,d,e
g: HypVINN (Ours)	0.7936 (0.0956)	b,c,d,e	0.9068 (0.0743)	c,d,e,f	1.1207 (0.5563)	c,d,e

The proposed hetero-modal *HypVINN* performs as well as the modality-specific models. Furthermore, *HypVINN* with multi-modal and standalone T1w input outperforms the 3D-UNet proposed by Billot, Bocchetta, et al. (2020)—the only other contemporary method for hypothalamus parcellation. Note: the statistical significance column (Signif.) indicates which other models the model outperforms (paired Wilcoxon signed-rank test, corrected  $p < 0.05$ ).

( $n = 44$ ). Finally, performance is assessed on the unseen test-set by the three metrics (Dice, HD95, and VS).

### 3.1.1. Comparison with the state-of-the-art

In Table 2, we present the similarity scores for the global segmentation performance of all evaluation metrics as well as significance indicators (corrected  $p < 0.05$ ). Here, we observe that *HypVINN* performs as well as the modality-specific models. In the T1w-only inference scenario, the *T1-VINN* outperforms *HypVINN* in Dice and HD95; however, there is no statistical difference between them. On the other hand, when T1w and T2w are available, *HypVINN* outperforms the multi-modal model in all evaluation metrics with statistical significance in VS. Furthermore, inputting only a T2w yields the lowest segmentation results from all benchmark models, and the T2w specialized network outranks the *HypVINN* with statistical significance. Additionally, we observe that for *HypVINN* the inclusion of both modalities improves segmentation performance compared to its single modality counterparts with statistical significance for all metrics in T2w and for T1w only in Dice. For the modality-specific models, *MM-VINN* and *T1-VINN* perform equally well with no statistical significance between them. Finally, our models (both T1 and multi-modal variants) outperform the T1 *3D-UNet* in our segmentation task with statistical significance.

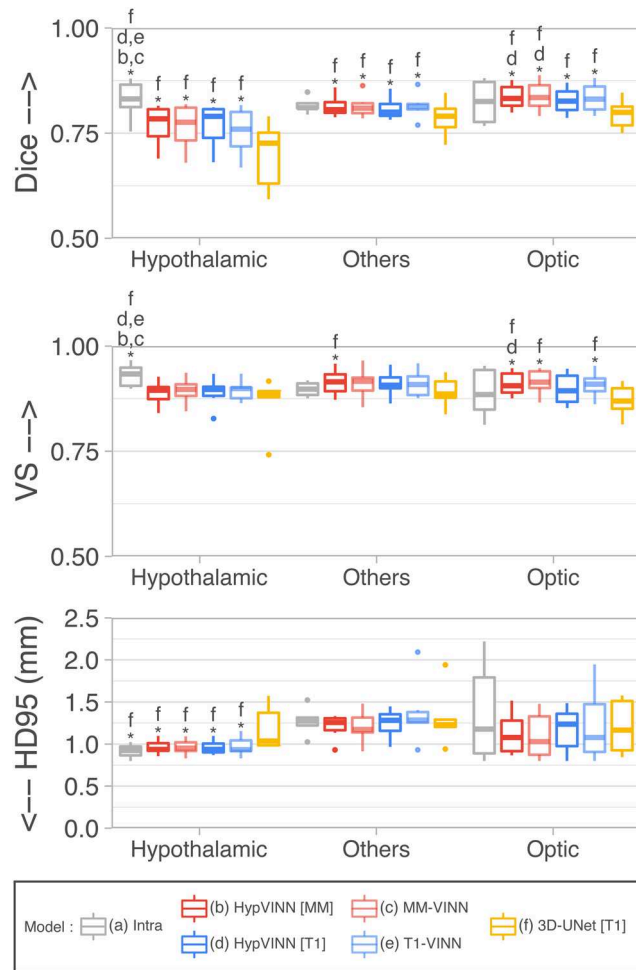
We additionally observe that the global results are not driven by any particular structure, as the per-structure results from *HypVINN* and the comparison models align

with their respective global outcomes. Furthermore, using a T2w scan as the only source for inferring information is consistently underperforming, at both the global and per-structure levels. For detailed per-structure performance results, see Appendix Figure A1.

Moreover, the contribution of T2-derived features can also be visualized in *HypVINN*'s learned global fusion weights where the T2-block weight ( $\approx 0.25$ ) has a much lower value than the T1-block weight ( $\approx 0.75$ ) starting already in early stages of training in all implemented networks as shown in Appendix Figure A2. Thus, performance is mainly driven by the T1-derived information, with T2w being only a support modality. For this reason, in the remaining experiments, we only use a T2w image in combination with a T1w image and not as a standalone modality.

### 3.1.2. Intra-rater reproducibility

In this section, we compare the performance of the automated methods against our main rater variability (i.e., intra-rater variability). The intra-rater variability puts the accuracy results into context, where it can be seen as the ideal automated method performance. We assess this variability by computing the similarity between the two sets of manual segmentations of the main rater in the in-house test-set. Note, in contrast to Section 3.1, all models are retrained on the full training dataset. It is important to note that the testing-set is still unseen for these models and is only used for final performance. These “final”



**Fig. 3.** Segmentation performance comparison on the in-house test-set between manual intra-rater scores vs. our proposed *HypVINN* and benchmark F-CNNs. *HypVINN* (dark red and dark blue) produces comparable results to the manual intra-rater agreement (gray). Note: similarity scores are presented for the hypothalamic, others, and optic regions. Additionally, a letter directly on top of a box plot indicates which other models the model significantly outperforms (paired Wilcoxon signed-rank test, corrected  $p < 0.05$ ).

models are additionally used for the generalizability (Section 3.2), reliability (Section 3.3), and sensitivity (Section 3.4) analyses.

In Figure 3, we present box plots for the three accuracy metrics (Dice, VS, and HD95) in the test-set for the three major regions (hypothalamic, optic, and others, see Section 2.2). We observe that our main rater has an overall good intra-rater agreement between annotation sessions (Global  $\rightarrow$  Dice = 0.8210, VS = 0.9100, HD95 = 1.1277 mm). Furthermore, all automated 2.5D methods perform equally well (Global  $\rightarrow$  T1-VINN: Dice = 0.7869, VS = 0.9017,


HD95 = 1.1638 mm; MM-VINN: Dice = 0.7937, VS = 0.9036, HD95 = 1.0723 mm; *HypVINN* with T1 input: Dice = 0.7905, VS = 0.8980, HD95 = 1.1103 mm; *HypVINN* with MM input: Dice = 0.7950, VS = 0.9008, HD95 = 1.0857 mm). Additionally, the 3D-UNet presents the lowest segmentation performance from all final models (Global  $\rightarrow$  Dice = 0.7435, VS = 0.8763, HD95 = 1.2347 mm).

The intra-rater scores outperform all the implemented automated methods in Dice and VS, with significant statistical differences present in the hypothalamic region structures (corrected  $p < 0.01$ ). Moreover, the HD95 inter-rater hypothalamic region results are significantly better than the ones of the 3D model. On the other hand, MM-VINN and *HypVINN* outperform the intra-rater results in recognizing tissue boundaries (HD95), even if no statistical significance can be inferred from the statistical test. We additionally observe that manually replicating the boundary outline in the structures from the others and optic regions is more challenging. Furthermore, we visually notice that all automated methods generate similar predictions to the manual ones, with the most considerable discrepancies in identifying the hypothalamus contour (outside boundaries), as illustrated in Figure 4. Moreover, the 3D model generates the noisiest hypothalamic edges from all implemented methods.

Finally, when comparing accuracy results between 2.5D automated methods, statistically significant differences are only present in Dice and VS for the optic region between *HypVINN* inference setups (corrected  $p < 0.05$ ) with the multi-modal input variation having a better performance (Dice: **0.8329** vs. 0.8238 and VS: **0.9119** vs. 9021). Nonetheless, we also observe improvements without statistical significance in hypothalamic region localization (Dice) and boundary detection (HD95) in structures from the others and optic regions. These results follow the previous section (Section 3.1.1), where *HypVINN* shows better segmentation results when all modalities are available. Moreover, the T1 and multi-modal 2.5D counterparts outperform the 3D model, aligning with previous findings.

### 3.2. Generalizability

In this section, we evaluate the robustness of the proposed hetero-modal model (*HypVINN*) to generalize to brain MRI scans with a different image resolution (1.0 mm isotropic) than the training one (0.8 mm isotropic). For this purpose, we utilize the MRI scans from the Rhineland Study (RS) in-house test-set ( $n = 6$ ) and a random subset ( $n = 9$ ) of the UK Biobank (UKB) dataset



**Fig. 4.** Comparison of the ground truth vs. predictions from the proposed *HypVINN* and comparison baselines for four participants of the in-house test-set. (A-D) All automated methods generate similar segmentation to the manual ones. However, differences are observed in the delineation of the hypothalamic contour. Furthermore, the *3D-UNet* presents the least smooth transitions between hypothalamic structures from all automated methods (red arrows). Note: each row represents a different participant with corresponding MRI modalities (T1-weighted (T1w) and T2w-weighted (T2w)), manual ground truth (GT), and automated generated segmentations on the coronal view. The color scheme for the visible structures is presented on the right.

that is manually annotated (see [Section 2.1](#)). For the Rhineland Study, as the MR scans and respective ground truth are at 0.8 mm isotropic resolution, we down-sample the pre-registered T1w and T2w scans from their native resolution to the desired 1.0 mm isotropic resolution. After the 1 mm scans are processed by the segmentation model, the resulting probability maps (i.e., soft-labels) are up-sampled to the original 0.8 mm resolution. Thereafter, hard labels are generated. This strategy prevents the down-sampling of manual labels to 1.0 mm, which introduces interpolation artefacts that could potentially decrease accuracy along boundaries, thereby impacting the analysis. On the other hand, no re-sampling is needed for the UK Biobank scans as this dataset is acquired and labeled at 1.0 mm resolution. However, multi-modal evaluation is not done for this dataset as T2w scans are not available. Therefore, we limit the generalizability analysis in the UK Biobank dataset to the performance of the standalone T1w input models. Finally, generalizability performance is assessed by the three similarity metrics (Dice, HD95, and VS) at the native resolution of the corresponding manual reference, except for volume similarity (VS) in the 1.0 mm Rhineland Study predictions. VS does not require spatial overlap between label

maps; thus, it can be computed without the need for re-sampling to the same resolution.

[Henschel et al. \(2022\)](#) demonstrated generalizability of *VINN*, *HM-VINN*'s parent architecture, to unseen resolutions. Their results, however, were achieved training with multi-resolution data, which is a different scenario to ours, where only 0.8 mm data is available. Therefore, here we further compare generalizability of our *HM-VINN* architecture to segment 1.0 mm MR scans to F-CNNs without resolution-independence mechanisms (*HM-CNN*). In *HM-CNN*, we replace the flexible network-integrated resolution-normalization step inside *HM-VINN* with a fixed scale transition. Furthermore, to isolate the contributions of the proposed resolution-independence scheme, we train both *HM-VINN* and *HM-CNN* with and without external scale augmentation (exSA). It is important to note that the *HM-VINN* + exSA (proposed *HypVINN*) used in this analysis is the one trained in [Section 3.1.2](#). Therefore, to ensure a fair comparison, all benchmarked networks are trained using the same procedure. We limited this analysis to only T1 input models as T1 is the primary MRI sequence for our segmentation task. Finally, in order to validate the robustness of *HypVINN* in both inference scenarios, we compare our method against the modality-specific



model from the previous section (i.e., *T1-VINN*, *MM-VINN*, and *3D-UNet*).

In [Figures 5](#) and [6](#), we present the generalizability results for the segmentation evaluation metrics in the hypothalamic, optic, and others regions for both datasets. For the first comparison analysis ([Fig. 5](#)), the inclusion during training of exSA in both *HM-VINN* (proposed *HypVINN*, [Fig. 5](#) blue) and *HM-CNN* ([Fig. 5](#) green) architectures shows better segmentation performance compared to their respective comparative baseline without exSA ([Fig. 5](#) orange and purple). Furthermore, we observe that the proposed *HypVINN* (*HM-VINN* +exSA) yields the best segmentation scores among all benchmark networks across different regions and metrics for both datasets, except for HD95 in the optic and hypothalamic structures for UKB. However, the differences in HD95 performance between our *HypVINN* and the *HM-VINN* (optic region) and *HM-CNN* +exSA (hypothalamic region) baselines in the UKB dataset are not statistically significant (corrected  $p > 0.1$ ). Lastly, as expected, the vanilla *HM-CNN* (no exSA or resolution-independence) fails in both datasets for all regions, showcasing the expected generalizability issues of a standalone F-CNN to out-of-distribution resolutions.

Analyzing the generalizability results between input modalities, we observed that even though models have not been trained at 1.0 mm resolution, they can generalize remarkably well, as illustrated in [Figures 6](#) and [7](#). For RS, no significant differences are found between 2.5D models except for the optic area where both multi-modal models outperform the T1-input *HypVINN* with statistical significance (corrected  $p < 0.02$ ; metric significance: Dice and VS both methods, and HD95 only *MM-VINN*). In UKB scans, the T1-input *HypVINN* outperforms the T1-specialized model in all metrics for the hypothalamic region. On the other hand, *T1-VINN* outranks our hetero-modal model in the others and optic regions. However, none of the above differences are statistically significant (corrected  $p > 0.1$ ). Finally, when comparing against the *3D-UNet* (which has been trained with external scale augmentation), the 2.5D models show in RS significantly better Dice scores for the hypothalamic and optic regions (corrected  $p < 0.02$ ). For UKB, the 2.5D models significantly outperform the *3D-UNet* in Dice and HD95 for the hypothalamic and others regions (corrected  $p < 0.01$ ).

### 3.3. Test-retest reliability

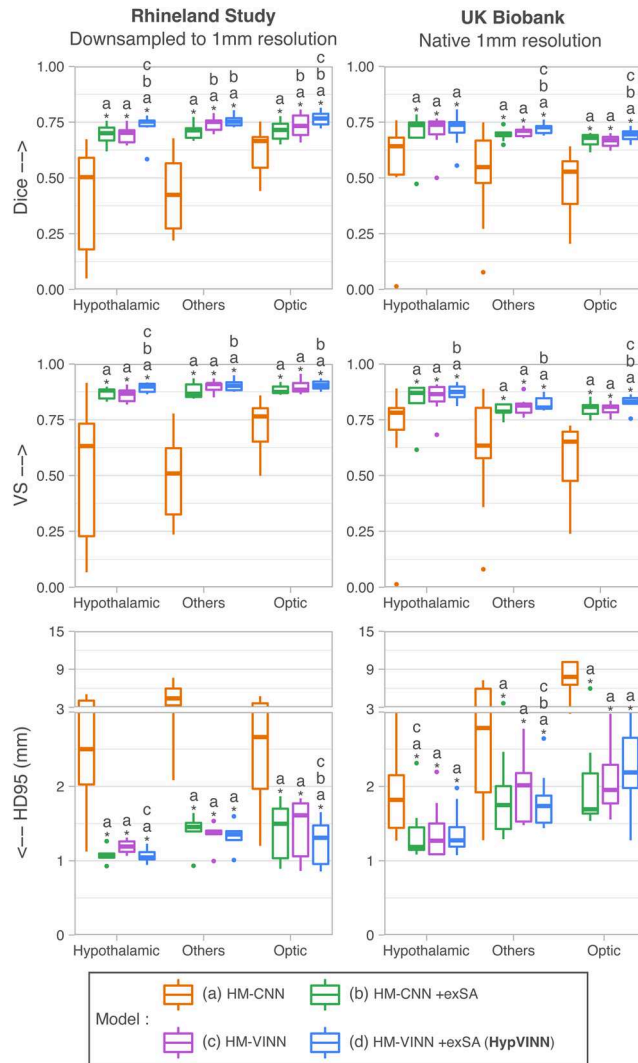
Assuming that brain anatomy does not change within the same MR session, a reliable method should generate the same (or very similar) volume estimates from repeated

in-session scans acquired under the same conditions (e.g., machine, acquisition protocol, region of interest). Here, we benchmark and evaluate the reliability of our proposed hetero-modal F-CNN to predict hypothalamic sub-regions and adjacent structure volumes in a test-retest scenario. For this purpose, we process the T1w and T2w scans from the test-retest dataset ( $n = 21$ ) not only with *HypVINN* but also with the benchmark models used in the previous sections (see [Sections 3.1.2](#) and [3.2](#)) except for the *3D-UNet* as it is the model with the lowest segmentation accuracy results. Since the test-retest dataset includes two T1w scans per participant and only a single T2w scan, the T2w is independently registered two times, each time using a different T1w as reference. Afterwards, we assess the reliability of the methods by computing volume similarity (VS) and intra-class correlation (ICC) between volume predictions across sequences. Finally, we compare the methods' volume similarity performance with a paired two-sided Wilcoxon signed-rank test.

All methods have an excellent agreement (ICC(A,1)  $> 0.95$ ) between volume predictions across sequences for all regions, as can be seen in [Appendix Table A5](#). Furthermore, all implemented methods perform equally well for VS in all regions (VS  $> 0.98$ ). Finally, we observe a statistically significant difference in the structures from the others region between *HypVINN* with multi-modal input and *T1-VINN* (VS: **0.9960** vs. 9927, corrected  $p < 0.05$ ).

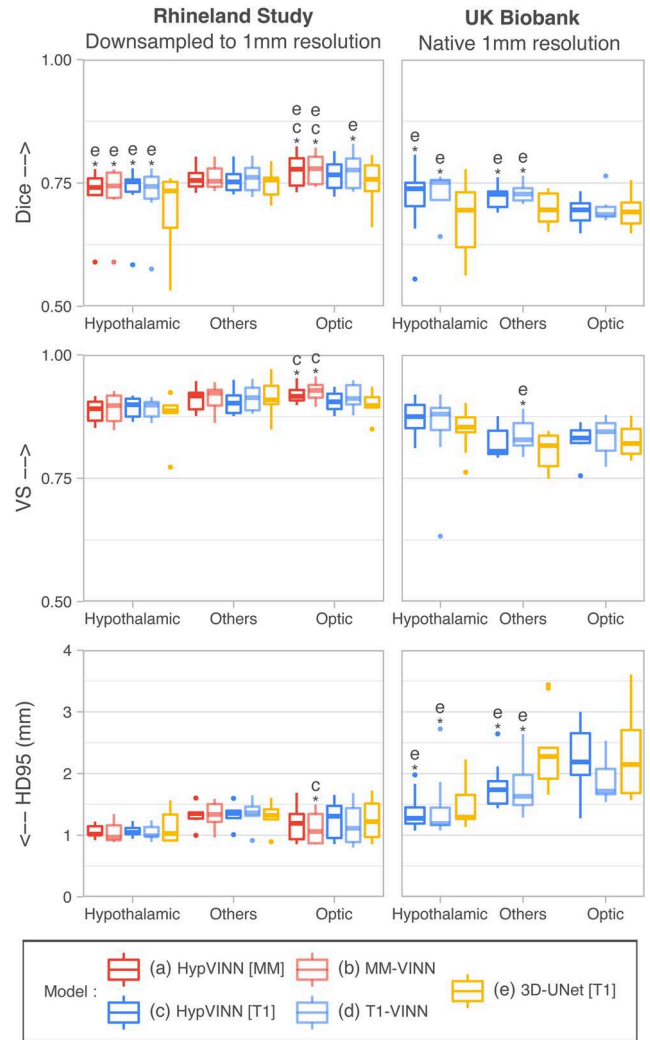
### 3.4. Sensitivity to age and sex effects

Previous studies have shown that men have a larger hypothalamus volume than women not only at a global level ([Isiklar et al., 2022](#)) but also at a sub-unit level ([Makris et al., 2013](#); [Thomas et al., 2019](#)). Therefore, in this section, we aim to use the automated hypothalamic volume estimates to replicate these findings and explore volume-age correlations in a general population, representing a feasible scenario in which our method will be used as the post-processing analysis pipeline. To this end, we process the T1w scans from the Rhineland Study ( $n = 463$ ) and UK Biobank ( $n = 535$ ) case-study datasets (see [Section 2.1](#)) with our proposed *HypVINN*. To further evaluate the robustness of our hetero-modal model to handle different modalities, we also assess the effects in the Rhineland cases when both pre-registered T1w and T2w scans are available at inference. Ideally, the direction of the effects should not be modified by the input scenarios (only T1w or T1w & T2w). We note that joint T1w & T2w analysis in the UK Biobank is not possible due to the absence of T2w scans.



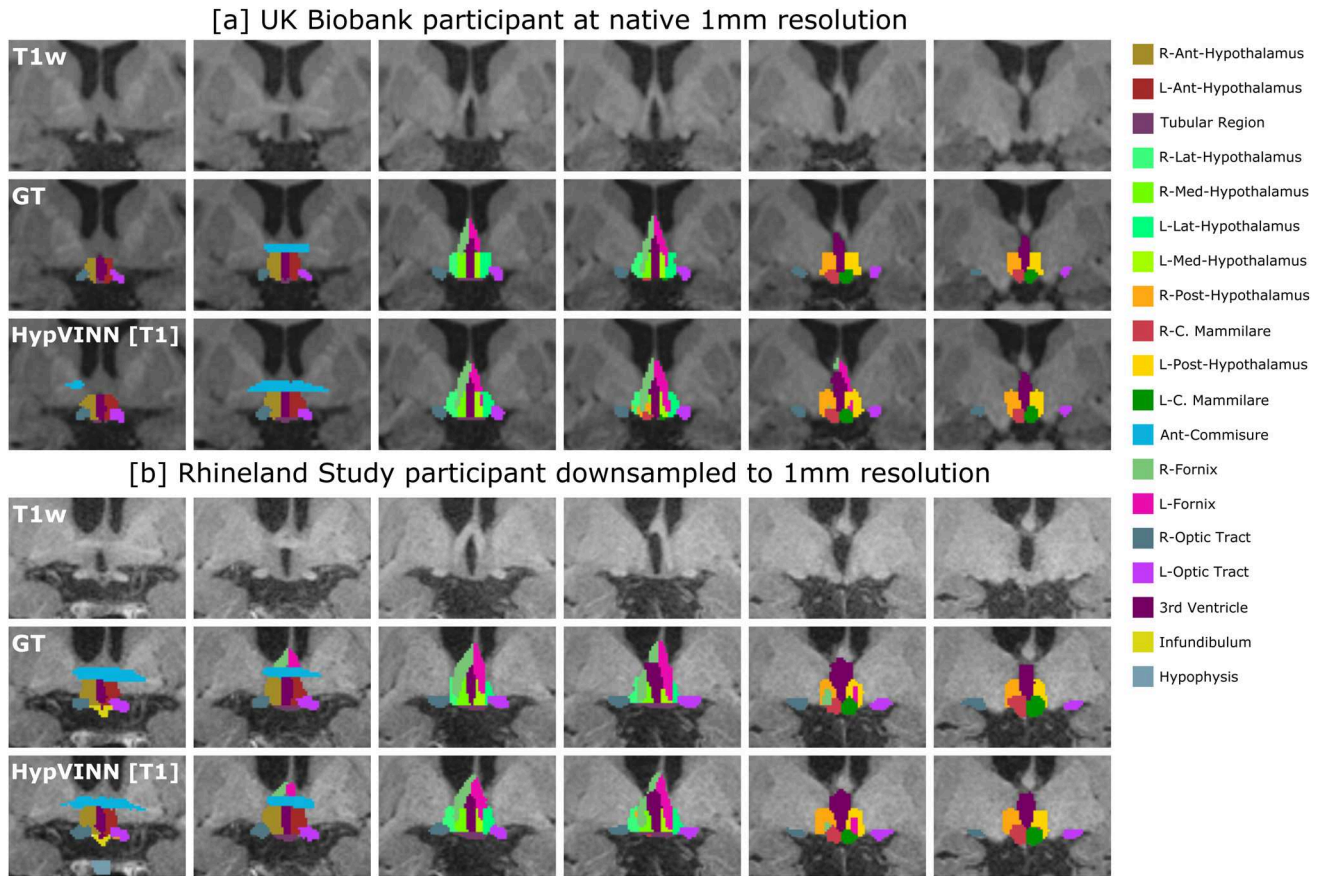
**Fig. 5.** Retrospectively benchmarking of single resolution (0.8 mm) trained networks to segment 1.0 mm T1w MR scans from the Rhineland Study and UK Biobank. Our proposed approach (*HypVINN*) consisting of the *HM-VINN* architecture plus external scale augmentation (+exSA, blue) outperforms other comparison baselines in both manually labeled datasets. Note: similarity scores are presented for the hypothalamic, others, and optic regions. Additionally, a letter directly on top of a box plot indicates which other models the model significantly outperforms (paired Wilcoxon signed-rank test, corrected  $p < 0.05$ ).

All generated predictions are visually inspected by an experienced rater. A total of 6 participants (1.29%) from the Rhineland Study (RS) and 15 participants (2.80%) from the UK Biobank (UKB) are excluded from the analysis sample due to segmentation errors (see, e.g., [Appendix Fig. A3](#)). For the remaining participants (RS:  $n = 457$ , UKB:  $n = 520$ ), bias field correction is performed for all T1w and



**Fig. 6.** Segmentation performance comparison between our proposed *HypVINN*, with multi-modal input (MM) and uni-modal T1 input (T1), vs. modality-specific models for segmenting 1.0 mm MR scans from the Rhineland Study and UK Biobank. *HypVINN* (dark red and dark blue) can generalize remarkably well to 1.0 mm MR scans independent of the provided MRI input. Note: similarity scores are presented for the hypothalamic, others, and optic regions. Additionally, a letter directly on top of a box plot indicates which other models the model significantly outperforms (paired Wilcoxon signed-rank test, corrected  $p < 0.05$ ).

T2w scans as a pre-processing step, and structure volume estimates are compensated for partial volume effects using *FastSurfer's* optimized Python re-implementation of *FreeSurfer's* `mri_segstats` command (`segstats.py`). Finally, for the total hypothalamus as well as for each of the hypothalamic sub-regions, we calculate the association per dataset of age and sex with the respective volumes using



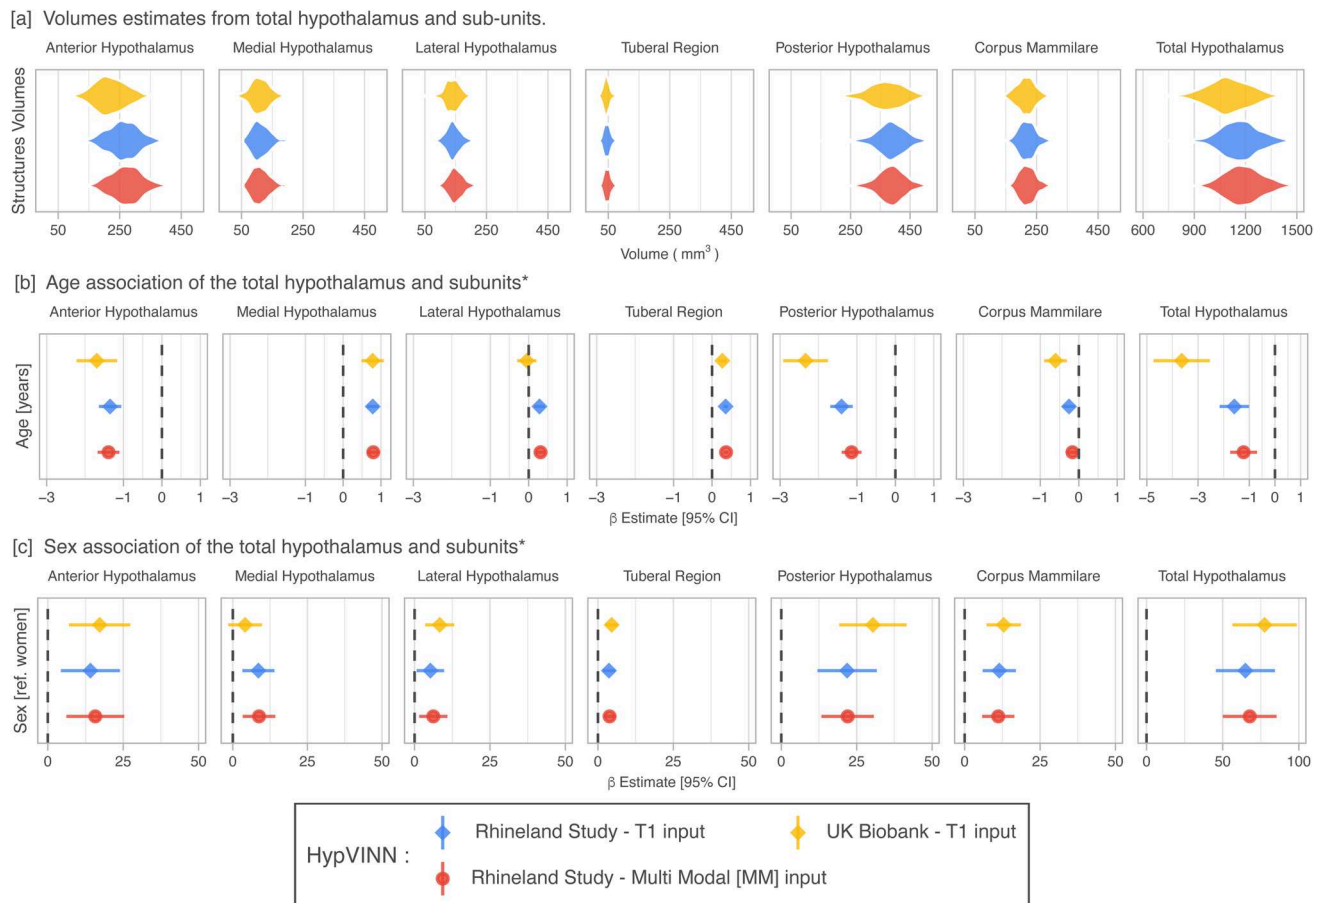
**Fig. 7.** Segmentation examples on the coronal view from our proposed *HypVINN* with T1 input and manual ground truth (GT) for one labeled 1.0 mm scan from the UKBiobank (a) and one 1.0 mm scan from the Rhineland Study unseen test-set (b). Even though our proposed method is not trained with 1.0 mm scans, it can generate accurate predictions at this resolution. Note: the color scheme for the visible structures is presented on the right.

independent multi-variable linear regression models. All models are further adjusted for head-size (estimated total intracranial volume,  $eTIV$ ), and RS models are also corrected for the T1w sequence version ( $T1_{seq}$ ) and T2w sequence version ( $T2_{seq}$ ). Furthermore, de-meaned versions of age ( $\hat{age}$ ) and  $eTIV$  ( $e\hat{TIV}$ ) are used in the association analysis (UKB-Model:  $Volume \sim \hat{age} + sex + e\hat{TIV}$ , RS-Model:  $Volume \sim \hat{age} + sex + e\hat{TIV} + T1_{seq} + T2_{seq}$ ). All statistical analyses are performed in R (R Core Team, 2020), and  $eTIV$  estimations are computed using *FreeSurfer* (Buckner et al., 2004). It is important to note that automated segmentations can be carried out without needing bias field corrected scans. Here, we correct the bias field in a pre-processing step primarily for the partial volume estimation, which is a post-processing step to the segmentation.

The predicted volumes for the total hypothalamus follow the results from smaller studies (Bocchetta et al., 2015; Chen et al., 2019; Makris et al., 2013; Rodrigues et al.,

2022; Schindler et al., 2013) with a similar global anatomical definition (from  $910 \text{ mm}^3$  to  $1580 \text{ mm}^3$ ) as can be seen in Figure 8a. For the sub-regions, we observe that the tubular region is the smallest segmented hypothalamic structure ( $\pm 45.9 \text{ mm}^3$ ) and the posterior hypothalamus the biggest one ( $\pm 379.3 \text{ mm}^3$ ). However, a direct comparison in size of our hypothalamic sub-regions with other studies is not possible due to the different segmentation protocols.

For both RS and UKB subsets, the total hypothalamus volumes significantly decreased ( $p < 0.001$ ) with age (see Fig. 8b). This negative association is also observed in the sub-regions except for the middle structures (e.g., tubular-region, medial and lateral hypothalamus), where the volumes are positively correlated with age. However, this positive correlation in all middle structures is not observed in the UKB, where a significant increase is not found for the lateral hypothalamus. Furthermore, all structures independent of the dataset, except for the medial hypothalamus in UKB, show statistically significant sex



**Fig. 8.** Hypothalamic volumes estimates (a) and volume associations with age (b) and sex (c) in participants from the Rhineland Study ( $n = 457$ ) and UK Biobank ( $N = 520$ ) for *HypVINN*. Age and sex effects on hypothalamic volume estimates in the Rhineland Study from *HypVINN*, independent of the provided MRI input, follow the same direction trend. Furthermore, our model replicates previous sex findings in both datasets corroborating the stability and sensitivity of our method. Note: Effects are obtained after accounting for head-size (eTIV) and modality sequence (only Rhineland Study).

differences ( $p < 0.05$ ) even after correcting for head-size, with men having larger hypothalamic volumes than women (see Fig. 8 c). These results are in line with previous findings (Isiklar et al., 2022; Makris et al., 2013; Thomas et al., 2019). Moreover, as expected, all inferred volumes are positively associated with eTIV ( $p < 0.01$ ).

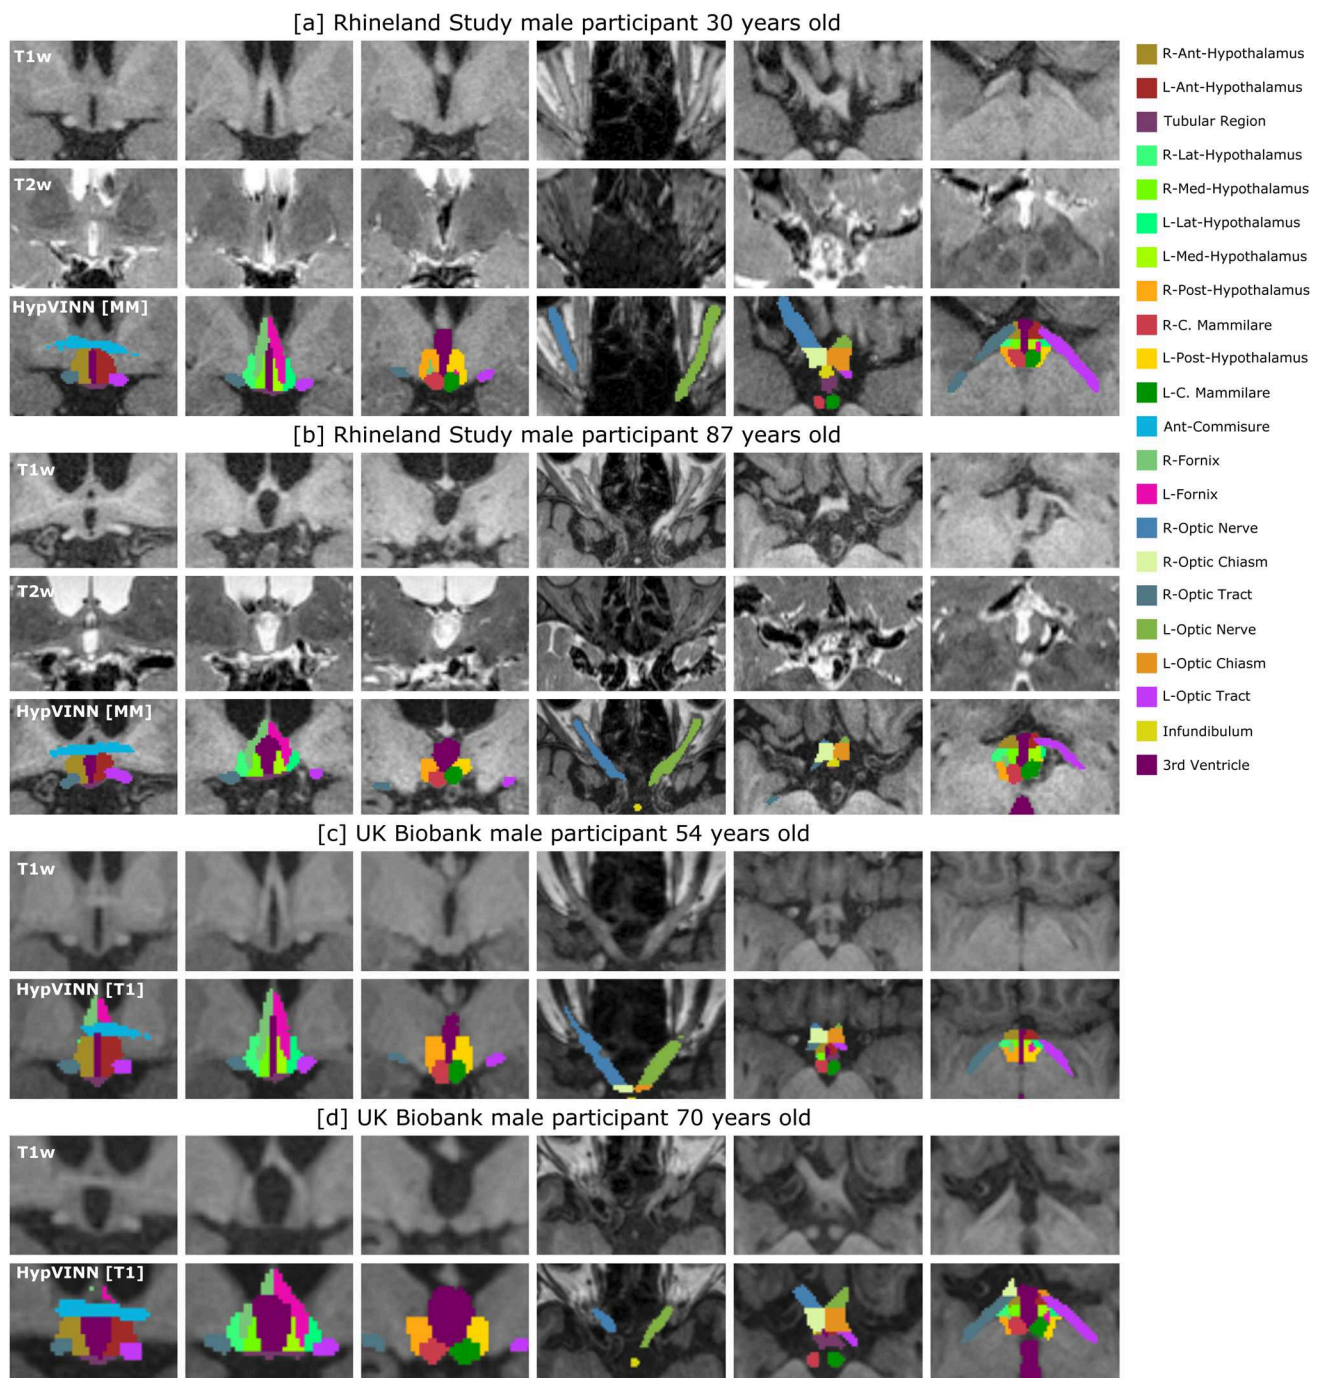
Independent of the provided MRI input, age and sex effects on hypothalamic volume estimates in the Rhineland Study using our *HypVINN* exhibit the same directional trends. Moreover, even though *HypVINN* is trained with all RS sequence versions, we observe differences between sequences; however, none of them are significant ( $p > 0.05$ ). Nevertheless, controlling for MRI sequences in any downstream statistical analysis is recommended when including image biomarkers obtained from multiple MRI sequences.

From the visual quality assessment, we observe that our tool performed very well in two different datasets; examples of correct segmentations for four random male participants with different ages can be observed in Figure 9. For the failing cases, we note that segmentation errors are mainly present when there is an unclear boundary of the hypothalamus due to severe enlargements of the third ventricle as illustrated in Appendix Figure A3.

#### 4. DISCUSSION

In this paper, we present the first hetero-modal model for automated sub-segmentation of the hypothalamus and adjacent structures on T1w and T2w brain MRI at isotropic 0.8 mm or 1 mm resolutions. The proposed model can generate accurate segmentations of the 24 different





**Fig. 9.** Examples of correct predictions in the Rhineland Study (a-b) and UK Biobank (c-d) from our proposed *HypVINN* with multi-modal [MM] or T1w only [T1] input for four unseen random male participants with different ages. Note: for each participant, T1w, T2w (only Rhineland Study participants), and *HypVINN* outcomes are presented. Furthermore, in each participant's row, the first three images display the different hypothalamic structures on the coronal view, and the remaining three images show all remaining structures on the axial view. The color lookup table for all visible structures is presented on the right.

structures in less than a minute from a standalone T1w image or by including an additional co-registered T2w image, without requiring multiple input-specific models, thus providing a robust, quick, and reliable solution for assessing hypothalamic volumes in small and large cohorts.

Firstly, we introduce a different segmentation protocol of the hypothalamus compared to the one proposed by Makris et al. (2013). Therefore, we re-train the only other contemporary method for hypothalamus sub-segmentation of 1 mm T1w images (Billot, Bocchetta, et al., 2020). The parcellation method of Makris et al. was developed for in-vivo semi-automatic hypothalamic segmentation using 1.5 T isotropic 1 mm MR images and was therefore necessarily less detailed than the one presented in this work. In general, we define the boundaries of the hypothalamus as a whole according to the same anatomical definitions and landmarks used by them. Yet, for sub-segmentation of the different hypothalamic sub-regions, we use a more fine-grained approach to take optimal advantage of the higher spatial resolution offered by the available 3 T 0.8 mm isotropic MR images. Consequently, our approach results in the sub-segmentation of more hypothalamic structures as detailed in Table 1. For example, whereas both the posterior hypothalamus and mammillary bodies were included under the label “posterior hypothalamus” in the parcellation scheme of Makris et al., our method provides separate volumetric estimates for each of these structures, which is of clinical relevance given that these structures operate in a functionally independent manner. Another noteworthy difference between the two parcellation schemes concerns the subdivision of the medial part of the hypothalamus: in contrast to Makris et al. who subdivided this region into a superior and an inferior tuberal region, we follow the more conventional neuroanatomical subdivision of this region into the medial and the lateral hypothalamus—using the fornix as the boundary between these two structures—and tubular region. For the tubular region, we group the tuberomammillary region, the median eminence, and the arcuate nucleus. Again, we opt for this approach to gain more detailed anatomical information about the various sub-structures of the hypothalamus. In addition, our method also provides automatic segmentation of several other important structures in the vicinity of the hypothalamus, for which, until now, no automated segmentation procedure has been available. Notably, these adjacent hypothalamic structures include the hypophysis (i.e., the pituitary gland), which is the body’s principal and most versatile endocrine gland responsible for the central

regulation of most other endocrine tissues throughout the body; the epiphysis, the site where the “sleep hormone” melatonin is synthesized; as well as all major structures of the central optic system, including the optic nerves, the optic chiasm, and the optic tracts.

Despite the small size of different sub-structures and low contrast on MR images, our novel deep-learning technique (*HypVINN*) can accurately segment all 24 structures even when input modalities are missing at inference time. *HypVINN* performs as well as state-of-the-art modality-specific F-CNNs. Passing a T2w scan as standalone input to *HypVINN* or to a specialized T2w model generates the lowest performance from all input variations (see Section 3.1.1). For our hetero-modal model, the difference in contribution between T1- and T2-derived information is quantifiable in the modality weights from the fusion module, with the weight of the T1-block ( $W_{T1}$ ) tripling the T2 one. Thus, an available T1w scan is more important for the current segmentation task than a T2w scan. Nonetheless, we demonstrate that including a T2 can still be beneficial for some structures as models with multi-modal information yield generally better segmentation performance.

Unequal performance between inference setups (i.e., available input modalities) was also reported in other hetero-modal deep-learning segmentation tasks, with higher results achieved when the primary modality was available (Dorent et al., 2019; Havaei et al., 2016; Varsavsky et al., 2018). In our case, preference for the T1 modality could be explained by the inherent modality bias from the manual annotation process. Our labeling protocol is mainly performed on the T1w scans, and the T2w scans are only used as a support modality as most anatomical boundaries are visible in T1. Hence, evaluating segmentation performance with the current manual labels is not entirely neutral across the various inference configurations. A more fair evaluation will require training and validation using manual annotations explicitly tailored to a structure’s visible anatomical characteristics in each input combination. However, generating  $2^m - 1$  manual labels per participant, where  $m$  represents the number of modalities, is not feasible as creating manual annotations for a single configuration is already expensive and time-consuming. Therefore, based on our findings, we recommend utilizing a T2w scan accompanied by a T1w scan (i.e., multi-modal input) and not as a standalone input for the current segmentation task.

Our hetero-modal model, when including a T1w image, exhibits segmentation performance in the range of the main rater variability (see Section 3.1.2). The

intra-rater variability can be seen as the ideal performance of the automated method as we use manually annotated labels from the main rater to train our F-CNNs. Therefore, it is challenging for an automated approach to outperform the intra-rater scores. Considering this, the accuracy in the hypothalamic region of our hetero-modal model and all benchmark methods is lower than the intra-rater agreement on all evaluation metrics. Yet, the underperformance in this region can also be attributed to the low MR contrast between neighboring structures, especially for the medial and lateral hypothalamus. Nonetheless, the segmentation results are en-par with other deep-learning techniques on similar brain segmentation tasks (i.e., small size and low contrast across anatomical boundaries) (Billot, Bocchetta, et al., 2020; Estrada et al., 2021).

*HypVINN* not only performs well on segmenting isotropic 0.8 mm T1w and T2w MR scans, but it also exhibits generalizability to isotropic 1 mm MR scans from the Rhineland Study and UK Biobank dataset (see Section 3.2). We demonstrate that utilizing the resolution-independence mechanism performs as well as external scale augmentations to handle unseen resolution when training with a single (0.8 mm) resolution. Furthermore, we show that resolution-independence combined with external scale augmentations (proposed) outperforms all other comparative baselines.

Furthermore, *HypVINN* performs equally well as modality-specific models in both 1 mm datasets. As expected, performance on the Rhineland Study data is higher than on the UK Biobank. The UK Biobank dataset consists of scans from a different cohort and is acquired with a different MRI acquisition protocol. Due to these dissimilarities, segmentation performance is not directly comparable. Nevertheless, the proposed *HypVINN* generalizes quite well to this external dataset. Finally, even though our model supports both 0.8 mm and 1 mm resolutions, we recommend to process 0.8 mm MR scans at their native resolution to obtain more detailed and precise predictions by leveraging the additional information present in the higher resolution. Note, our proposed model also shows promising results in the high-resolutional MRI scans from the Human Connectome Project (HCP) young adult and lifespan pilot project datasets (Bookheimer et al., 2019; Harms et al., 2018; Van Essen et al., 2012); see Appendix Figure A4 for prediction examples of our tool in HCP scans.

Throughout this work, we compare our *HypVINN* against the re-trained version of the *3D-UNet* with extensive data augmentations proposed by Billot, Bocchetta, et al. (2020a) for hypothalamus sub-segmentation. Our

results demonstrate that our method not only outperforms the *3D-UNet* in terms of segmentation accuracy (see Sections 3.1.1 and 3.1.2) but also exhibits better generalizability across both comparative datasets (see Section 3.2). Additionally, the training process for the *3D-UNet* using the authors' released implementation and recommended training parameters takes approximately 100 hours per model using the GPU setup described in Section 2.3.3. In contrast, back-to-back training of the three F-CNNs that compose our *HypVINN* takes around 19 hours (roughly 6 hours per F-CNN). Therefore, besides outperforming the contemporary method, our approach can be (re)trained more efficiently with a lower carbon footprint.

As demonstrated in the Rhineland Study data, all automated methods exhibit excellent test-retest agreement between in-session volume estimates (see Section 3.3). Additionally, our *HypVINN* shows high robustness and generalizability across the general population of the Rhineland Study and UK Biobank case-study datasets, with only 21 cases (2.10%) between the two datasets being excluded from the age and sex analysis due to segmentation errors (see Section 3.4). The most common factor for our pipeline to fail is a severe deformation of the third ventricle (i.e., out-of-distribution cases), which generates unclear hypothalamic boundaries, as illustrated in Appendix Figure A3. Therefore, careful inspection is recommended when using our tool in aging populations and clinical cohorts, as the prevalence of large ventricles increases with age and certain diseases (e.g., Alzheimer's disease, Parkinson's disease, etc.). We recommend visually inspecting the predictions from scans with pathological changes and from volumetric outliers within the cohort before including them in any downstream analysis, particularly outliers from the third ventricle and medial/lateral hypothalamus. Although volumetric outlier detection can help identify predictions with significant failures, more robust quality control tools are desirable. However, developing these tools is outside this paper's scope and will be future work.

In line with previous studies on smaller datasets (Isiklar et al., 2022; Makris et al., 2013; Thomas et al., 2019), we also find that the volume of the total hypothalamus is larger in men compared to women. However, our analyses in two substantially larger population-based cohorts revealed that the volumes of virtually all hypothalamic substructures are significantly larger in men independent of head size. Our findings thus warrant further detailed association studies to investigate the clinical relevance of these pronounced sex differences in



the human hypothalamus. On the other hand, the derived age effects from small-scale studies present inconsistent results for the different hypothalamic substructures, except for the total hypothalamus whose total volume decreases with age (Billot, Bocchetta, et al., 2020; Bocchetta et al., 2015; Isiklar et al., 2022; Makris et al., 2013). Our method's total hypothalamic volume estimates also replicate this negative correlation with age. Furthermore, although most hypothalamic regions atrophy with increasing age, the volume of the middle/tuberal region of the hypothalamus significantly increases with age. This finding is novel and could imply that specific hypothalamus regions could be resistant to age-associated atrophy. Indeed, the paraventricular nucleus contained within the medial hypothalamic region exhibits a striking stability in terms of neuronal numbers, both with age and in the context of common neurodegenerative diseases such as Alzheimer's disease (Lucassen et al., 1994). These findings thus underscore the need for further large-scale studies into the differential effects of age on different hypothalamic substructures.

In conclusion, we demonstrate that *HypVINN* can successfully identify the desired structures with similar or better performance than state-of-the-art modality-specific models regarding segmentation accuracy, generalizability, and test-retest reliability. Furthermore, the fact that *HypVINN* replicates previous age and sex findings on large unseen subsets of the Rhineland Study and the UK Biobank corroborates the stability and sensitivity of our method. Moreover, our hypothalamic sub-segmentation tool generates accurate segmentations regardless of whether both T1w and T2w images are available or just a single T1w image. However, utilizing both modalities results in slightly improved segmentation outcomes.

Future work will focus on supporting a wider range of resolution by training our *HypVINN* with multi-resolution, thus fully exploiting the advantages of using a voxel-size independent F-CNN (VINN) (Henschel et al., 2022). Moreover, we will also focus on improving the robustness of our tool to out-of-distribution cases (e.g., severe deformation of the third ventricle). Since *HypVINN* is based on deep learning, boosting the robustness to these cases can potentially be achieved by retraining with manual annotations created on participants with low segmentation quality or by applying realistic non-linear deformations as an additional data augmentation during the training process (Faber et al., 2022). Finally, extending the input flexibility of our tool to scenarios where input scans are at different resolutions (mixed resolutions) is also of interest, as it could allow the deployment of our

tool in more scenarios where HighRes data are unavailable in all modalities.

Overall, we introduce *HypVINN*—the first hetero-modal deep-learning method for hypothalamic sub-segmentation and segmentation of other adjacent structures, such as the hypophysis, epiphysis, and major structures of the central optic system. The proposed method offers a more detailed parcellation of the hypothalamus compared to the only other contemporary automated method (Billot, Bocchetta, et al., 2020). Additionally, it can generate accurate segmentations from T1w and T2w MR images at isotropic 0.8 mm or 1 mm resolutions. Finally, *HypVINN* will be incorporated into the *FastSurfer* neuroimaging software suite, thus providing an easy-to-use alternative for more reliable assessment of hypothalamic imaging-derived phenotypes.

#### DATA AND CODE AVAILABILITY

This work uses MRI data from the Rhineland Study and UK Biobank. The Rhineland Study data are not publicly available because of data protection regulations. However, access can be provided to scientists in accordance with the Rhineland Study's Data Use and Access Policy. Requests to access the data should be directed to Dr. Monique Breteler at [RS-DUAC@dzne.de](mailto:RS-DUAC@dzne.de). UK Biobank data are available through a procedure described at <http://www.ukbiobank.ac.uk/using-the-resource/>.

The method presented in this article will be made publicly available on Github (<https://github.com/Deep-MI/FastSurfer>) upon acceptance.

#### AUTHOR CONTRIBUTIONS

Santiago Estrada: Methodology, Software, Validation, Formal analysis, Investigation, Conceptualization, Writing—original draft, Writing—review & editing, and Visualization. David Kügler: Conceptualization, Methodology, Validation, Data Curation, Writing—original draft, and Writing—review & editing. Emad Bahrami: Investigation, Software, Validation, Data Curation, and Writing—original draft. Peng Xu: Data Curation, Investigation, Validation, and Writing—original draft. Dilshad Mousa: Data Curation, Investigation. Monique M.B. Breteler: Supervision, Funding acquisition, Resources, and Writing—review & editing. N. Ahmad Aziz: Conceptualization, Validation, Resources, Writing—original draft, Writing—review & editing, Supervision, and Funding acquisition. Martin Reuter: Conceptualization, Validation, Resources, Writing—original draft, Writing—review & editing, Supervision, Project administration, and Funding acquisition.



## DECLARATION OF COMPETING INTEREST

The authors declare that they have no conflict of interest.

## ETHICS STATEMENT

This work uses MRI data from participants of the Rhineland Study and UK Biobank. Participants in both studies gave written informed consent in accordance with the ethical guidelines of the individual studies. The Rhineland study is carried out in accordance with the recommendations of the International Council for Harmonisation (ICH) Good Clinical Practice (GCP) standards (ICH-GCP). UK Biobank had obtained ethics approval from the North West Multicentre Research Ethics Committee.

## 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 of Germany (031L0206, 01GQ1801), the Chan Zuckerberg Initiative (Project Fast-Surfer, Grant Number: EOSS5 2022-252594), the Helmholtz-Al project DeGen (ZT-I-PF-5-078), an Alzheimer's Association Research Grant (Award Number: AARG-19-616534), and NIH (R01 LM012719, R01 AG064027, R56 MH121426, and P41 EB030006). Peng Xu was supported by a scholarship from China Scholarship Council, and N. Ahmad Aziz was supported by a European Research Council Starting Grant (Number: 101041677).

This research has been conducted using the UK Biobank Resource under Application Number 82056. Data in the appendix were also provided in part by the publicly available Human Connectome Project (HCP), 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.

## REFERENCES

- Ahmed, R. M., Halliday, G., & Hodges, J. R. (2021). Hypothalamic symptoms of frontotemporal dementia disorders. *Handbook of Clinical Neurology*, 182, 269–280. <https://doi.org/10.1016/b978-0-12-819973-2.00019-8>
- Alfaro-Almagro, F., Jenkinson, M., Bangerter, N. K., Andersson, J. L., Griffanti, L., Douaud, G., Sotiropoulos, S. N., Jbabdi, S., Hernandez-Fernandez, M., Vallee, E., Vidaurre, D., Webster, M., McCarthy, P., Rorden, C., Daducci, A., Alexander, D. C., Zhang, H., Dragonu, I., Matthews, P. M., Miller, K. L., & Smith, S. M. (2018). Image processing and quality control for the first 10,000 brain imaging datasets from UK biobank. *NeuroImage*, 166, 400–424. <https://doi.org/10.1016/j.neuroimage.2017.10.034>
- Avery, R. A., Mansoor, A., Idrees, R., Biggs, E., Alsharid, M. A., Packer, R. J., & Linguraru, M. G. (2016). Quantitative MRI criteria for optic pathway enlargement in neurofibromatosis type 1. *Neurology*, 86, 2264–2270. <https://doi.org/10.1212/wnl.0000000000002771>
- Baroncini, M., Jissendi, P., Balland, E., Besson, P., Pruvo, J.-P., Francke, J.-P., Dewailly, D., Blond, S., & Prevot, V. (2012). MRI atlas of the human hypothalamus. *NeuroImage*, 59, 168–180. <https://doi.org/10.1016/j.neuroimage.2011.07.013>
- Beliveau, V., Nørgaard, M., Birkel, C., Seppi, K., & Scherfler, C. (2021). Automated segmentation of deep brain nuclei using convolutional neural networks and susceptibility weighted imaging. *Human Brain Mapping*, 42, 4809–4822. <https://doi.org/10.1002/hbm.25604>
- 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. <https://doi.org/10.1016/j.neuroimage.2020.117287>
- Billot, B., Colin, Y., Magdamo Cheng, Das, S., & Iglesias, J. E. (2023). Robust machine learning segmentation for large-scale analysis of heterogeneous clinical brain MRI datasets. *Proceedings of the National Academy of Sciences (PNAS)*, 120, e2216399120. <https://doi.org/10.1073/pnas.2216399120>
- Billot, B., Greve, D. N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A. V., & Iglesias, J. E. (2023). Synthseg: Segmentation of brain MRI scans of any contrast and resolution without retraining. *Medical Image Analysis*, 86, 102789. <https://doi.org/10.1016/j.media.2023.102789>
- Billot, B., Greve, D. N., Van Leemput, K., Fischl, B., Iglesias, J. E., & Dalca, A. (2020). A learning strategy for contrast-agnostic MRI segmentation. In *Proceedings of the Third Conference on Medical Imaging with Deep Learning* (pp. 75–93). PMLR. <https://doi.org/10.1109/isbi48211.2021.9434113>
- Bocchetta, M., Gordon, E., Manning, E., Barnes, J., Cash, D. M., Espak, M., Thomas, D. L., Modat, M., Rossor, M. N., Warren, J. D., Ourselin, S., Frisoni, G. B., & Rohrer, J. D. (2015). Detailed volumetric analysis of the hypothalamus in behavioral variant frontotemporal dementia. *Journal of Neurology*, 262, 2635–2642. <https://doi.org/10.1007/s00415-015-7885-2>
- Bookheimer, S. Y., Salat, D. H., Terpstra, M., Ances, B. M., Barch, D. M., Buckner, R. L., Burgess, G. C., Curtiss, S. W., Diaz-Santos, M., Elam, J. S., Fischl, B., Greve, D. N., Hagy, H. A., Harms, M. P., Hatch, O. M., Hedden, T., Hodge, C., Japardi, K. C., Kuhn, T. P., ... Yacoub, E. (2019). The lifespan human connectome project in aging: An overview. *NeuroImage*, 185, 335–348. <https://doi.org/10.1016/j.neuroimage.2018.10.009>
- Brenner, D., Stirnberg, R., Pracht, E. D., & Stöcker, T. (2014). Two-dimensional accelerated MP-RAGE imaging with flexible linear reordering. *Magnetic Resonance Materials in Physics, Biology and Medicine*, 27, 455–462. <https://doi.org/10.1007/s10334-014-0430-y>
- 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, P92. <https://doi.org/10.1016/j.jalz.2014.05.172>
- Breuer, F. A., Blaimer, M., Mueller, M. F., Seiberlich, N., Heidemann, R. M., Griswold, M. A., & Jakob, P. M. (2006). Controlled aliasing in volumetric parallel imaging (2D CAIPIRINHA). *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 55, 549–556. <https://doi.org/10.1002/mrm.20787>
- 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, 724–738. <https://doi.org/10.1016/j.neuroimage.2004.06.018>
- 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. *Magnetic Resonance in Medicine*, 60, 640–649. <https://doi.org/10.1002/mrm.21680>
- Chaichana, K., & Quinones-Hinojosa, A. (2019). *Comprehensive overview of modern surgical approaches to intrinsic brain tumors*. Academic Press. <https://doi.org/10.1016/b978-0-12-811783-5.00026-4>
- Chen, Z., Chen, X., Liu, M., Ma, L., & Yu, S. (2019). Volume of hypothalamus as a diagnostic biomarker of chronic migraine. *Frontiers in Neurology*, 10, 606. <https://doi.org/10.3389/fneur.2019.00606>
- Di Martino, A., O'Connor, D., Chen, B., Alaerts, K., Anderson, J. S., Assaf, M., Balsters, J. H., Baxter, L., Beggiato, A., Bernaerts, S., Blanken, L. M. E., Bookheimer, S. Y., Braden, B. B., Byrge, L., Castellanos, F. X., Dapretto, M., Delorme, R., Fair, D. A., Fishman, I., ... Milham, M. P. (2017). Enhancing studies of the connectome in autism using the autism brain imaging data exchange II. *Scientific Data*, 4, 1–15. <https://doi.org/10.1038/sdata.2017.10>
- Dice, L. R. (1945). Measures of the amount of ecologic association between species. *Ecology*, 26, 297–302. <https://doi.org/10.2307/1932409>
- Dorent, R., Joutard, S., Modat, M., Ourselin, S., & Vercauteren, T. (2019). Hetero-modal variational encoder-decoder for joint modality completion and segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 74–82). Springer. [https://doi.org/10.1007/978-3-030-32245-8\\_9](https://doi.org/10.1007/978-3-030-32245-8_9)
- Dudás, B. (2021). Chapter 3 - Anatomy and cytoarchitectonics of the human hypothalamus. *Handbook of Clinical Neurology*, 179, 45–66. <https://www.sciencedirect.com/science/article/pii/B9780128199756000017>
- 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. *Magnetic Resonance in Medicine*, 83, 1471–1483. <https://doi.org/10.1002/mrm.28022>
- Estrada, S., Lu, R., Diers, K., Zeng, W., Ehses, P., Stöcker, T., Breteler, M. M., & Reuter, M. (2021). Automated olfactory bulb segmentation on high resolution t2-weighted MRI. *NeuroImage*, 242, 118464. <https://doi.org/10.1016/j.neuroimage.2021.118464>
- Faber, J., Kügler, D., Bahrami, E., Heinz, L.-S., Timmann, D., Ernst, T. M., Deike-Hofmann, K., Klockgether, T., van de Warrenburg, B., van Gaalen, J., Reetz, K., Romanzetti, S., Oz, G., Joers, J. M., Diedrichsen, J., & ESMI MRI Study Group. (2022). Cerebnet: A fast and reliable deep-learning pipeline for detailed cerebellum sub-segmentation. *NeuroImage*, 264, 119703. <https://doi.org/10.1016/j.neuroimage.2022.119703>
- Fischl, B. (2012). Freesurfer. *NeuroImage*, 62, 774–781. <https://doi.org/10.1016/j.neuroimage.2012.01.021>
- Fischl, B., Salat, D. H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., Van Der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., Montillo, A., Makris, N., Rosen, B., & Dale, A. M. (2002). Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain. *Neuron*, 33, 341–355. [https://doi.org/10.1016/s0896-6273\(02\)00569-x](https://doi.org/10.1016/s0896-6273(02)00569-x)
- Fronczek, R., Overeem, S., Lee, S. Y., Hegeman, I. M., Van Pelt, J., Van Duinen, S. G., Lammers, G. J., & Swaab, D. F. (2007). Hypocretin (orexin) loss in parkinson's disease. *Brain*, 130, 1577–1585. <https://doi.org/10.1093/brain/awm090>
- Glasser, M. F., Sotiropoulos, S. N., Wilson, J. A., Coalson, T. S., Fischl, B., Andersson, J. L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J. R., Van Essen, D. C., Jenkinson, M., & for the WU-Minn HCP Consortium. (2013). The minimal preprocessing pipelines for the human connectome project. *NeuroImage*, 80, 105–124. <https://doi.org/10.1016/j.neuroimage.2013.04.127>
- Greve, D. N., Billot, B., Cordero, D., Hoopes, A., Hoffmann, M., Dalca, A. V., Fischl, B., Iglesias, J. E., & Augustinack, J. C. (2021). A deep learning toolbox for automatic segmentation of subcortical limbic structures from MRI images. *NeuroImage*, 244, 118610. <https://doi.org/10.1016/j.neuroimage.2021.118610>
- 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, 1202–1210. <https://doi.org/10.1002/mrm.10171>
- Gulban, O. F., Nielson, D., Poldrack, R., John Lee, Gorgolewski, C., Vanessasaurus, & Ghosh, S. (2019). *poldracklab/pydeface: v2.0.0*. <https://doi.org/10.5281/zenodo.3524401>
- Güngör, A., Baydin, S., Middlebrooks, E. H., Tanriover, N., Isler, C., & Rhoton, A. L. (2017). The white matter tracts of the cerebrum in ventricular surgery and hydrocephalus. *Journal of Neurosurgery*, 126, 945–971. <https://doi.org/10.3171/2016.1.jns152082>
- Harms, M. P., Somerville, L. H., Ances, B. M., Andersson, J., Barch, D. M., Bastiani, M., Bookheimer, S. Y., Brown, T. B., Buckner, R. L., Burgess, G. C., Coalson, T. S., Chappell, M. A., Dapretto, M., Douaud, G., Fischl, B., Glasser, M. F., Greve, D. N., Hodge, C., Jamison, K. W., ... Yacoub, E. (2018). Extending the human connectome project across ages: Imaging protocols for the lifespan development and aging projects. *NeuroImage*, 183, 972–984. <https://doi.org/10.1016/j.neuroimage.2018.09.060>
- Havaei, M., Guizard, N., Chapados, N., & Bengio, Y. (2016). HeMIS: Hetero-modal image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention* (pp. 469–477). Springer. [https://doi.org/10.1007/978-3-319-46723-8\\_54](https://doi.org/10.1007/978-3-319-46723-8_54)

- 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. <https://doi.org/10.1016/j.neuroimage.2020.117012>
- Henschel, L., Kügler, D., & Reuter, M. (2022). FastSurferVINN: Building resolution-independence into deep learning segmentation methods—A solution for HighRes brain MRI. *NeuroImage*, 251, 118933. <https://doi.org/10.1016/j.neuroimage.2022.118933>
- Hofmann, M., Steinke, F., Scheel, V., Charpiat, G., Farquhar, J., Aschoff, P., Brady, M., Schölkopf, B., & Pichler, B. J. (2008). MRI-based attenuation correction for PET/MRI: A novel approach combining pattern recognition and atlas registration. *Journal of Nuclear Medicine*, 49, 1875–1883. <https://doi.org/10.2967/jnumed.107.049353>
- Huttenlocher, D. P., Klanderman, G. A., & Rucklidge, W. J. (1993). Comparing images using the hausdorff distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15, 850–863. <https://doi.org/10.1109/34.232073>
- Iglesias, J. E., Billot, B., Balbastre, Y., Tabari, A., Conklin, J., González, R. G., Alexander, D. C., Golland, P., Edlow, B. L., Fischl, B., & for the Alzheimer's Disease Neuroimaging Initiative. (2021). Joint super-resolution and synthesis of 1 mm isotropic MP-RAGE volumes from clinical MRI exams with scans of different orientation, resolution and contrast. *NeuroImage*, 237, 118206. <https://doi.org/10.1016/j.neuroimage.2021.118206>
- Isiklar, S., Turan Ozdemir, S., Ozkaya, G., & Ozpar, R. (2022). Hypothalamic volume and asymmetry in the pediatric population: A retrospective MRI study. *Brain Structure and Function*, 227, 2489–2501. <https://doi.org/10.1007/s00429-022-02542-6>
- 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. *Medical Image Analysis*, 36, 61–78. <https://doi.org/10.1016/j.media.2016.10.004>
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. ICLR. <https://arxiv.org/pdf/1412.6980.pdf>
- Lemaire, J.-J., Frew, A. J., McArthur, D., Gorgulho, A. A., Alger, J. R., Salomon, N., Chen, C., Behnke, E. J., & De Salles, A. A. (2011). White matter connectivity of human hypothalamus. *Brain Research*, 1371, 43–64. <https://doi.org/10.1016/j.brainres.2010.11.072>
- Liguori, C., Romigi, A., Nuccetelli, M., Zannino, S., Sancesario, G., Martorana, A., Albanese, M., Mercuri, N. B., Izzi, F., Bernardini, S., Nitti, A., Sancesario, G. M., Sica, F., Marciari, M. G., & Placidi, F. (2014). Orexinergic system dysregulation, sleep impairment, and cognitive decline in Alzheimer disease. *JAMA Neurology*, 71, 1498–1505. <https://doi.org/10.1001/jamaneurol.2014.2510>
- Loshchilov, I., & Hutter, F. (2019). Decoupled weight decay regularization. ICLR. <https://openreview.net/pdf?id=Bkg6RiCqY7>
- Lucassen, P., Salehi, A., Pool, C., Gonatas, N., & Swaab, D. (1994). Activation of vasopressin neurons in aging and Alzheimer's disease. *Journal of Neuroendocrinology*, 6, 673–679. <https://doi.org/10.1111/j.1365-2826.1994.tb00634.x>
- Makris, N., Swaab, D. F., van der Kouwe, A., Abbs, B., Boriel, D., Handa, R. J., Tobet, S., & Goldstein, J. M. (2013). Volumetric parcellation methodology of the human hypothalamus in neuroimaging: Normative data and sex differences. *NeuroImage*, 69, 1–10. <https://doi.org/10.1016/j.neuroimage.2012.12.008>
- McGraw, K. O., & Wong, S. P. (1996). Forming inferences about some intraclass correlation coefficients. *Psychological Methods*, 1, 30. <https://doi.org/10.1037/1082-989x.1.1.30>
- Merkel, D. (2014). Docker: Lightweight Linux containers for consistent development and deployment. *Linux Journal*, 2014, 2. <https://dl.acm.org/doi/10.5555/2600239.2600241>
- Milchenko, M., & Marcus, D. (2013). Obscuring surface anatomy in volumetric imaging data. *Neuroinformatics*, 11, 65–75. <https://doi.org/10.1007/s12021-012-9160-3>
- Miller, K. L., Alfaro-Almagro, F., Bangerter, N. K., Thomas, D. L., Yacoub, E., Xu, J., Bartsch, A. J., Jbabdi, S., Sotiropoulos, S. N., Andersson, J. L., Griffanti, L., Douaud, G., Okell, T. W., Weale, P., Dragonu, I., Garratt, S., Hudson, S., Collins, R., Jenkinson, M., Matthews, P. M., & Smith, S. M. (2016). Multimodal population brain imaging in the uk biobank prospective epidemiological study. *Nature Neuroscience*, 19, 1523–1536. <https://doi.org/10.1038/nn.4393>
- Milletari, F., Navab, N., & Ahmadi, S. (2016). V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In *Proceedings of the 2016 Fourth International Conference on 3D Vision (3DV)* (pp. 565–571). <https://doi.org/10.1109/3dv.2016.79>
- Mugler III, J. P. (2014). Optimized three-dimensional fast-spin-echo MRI. *Journal of Magnetic Resonance Imaging*, 39, 745–767. <https://doi.org/10.1002/jmri.24542>
- Orbes-Arteaga, M., Cárdenas-Peña, D., Álvarez, M. A., Orozco, A. A., & Castellanos-Dominguez, G. (2015). Magnetic resonance image selection for multi-atlas segmentation using mixture models. In *Iberoamerican Congress on Pattern Recognition* (pp. 391–399). Springer. [https://doi.org/10.1007/978-3-319-25751-8\\_47](https://doi.org/10.1007/978-3-319-25751-8_47)
- Park, J., Han, J. W., Suh, S. W., Byun, S., Han, J. H., Bae, J. B., Kim, J. H., & Kim, K. W. (2020). Pineal gland volume is associated with prevalent and incident isolated rapid eye movement sleep behavior disorder. *Aging (Albany, NY)*, 12, 884. <https://doi.org/10.18632/aging.102661>
- 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. In 31st Conference on Neural Information Processing Systems (NIPS 2017), December 4–9, 2017, Long Beach, CA, USA. [https://www.scirp.org/\(S\(351jmbntvnsjt1aadkozje\)\)/reference/referencespapers.aspx?referenceid=2530087](https://www.scirp.org/(S(351jmbntvnsjt1aadkozje))/reference/referencespapers.aspx?referenceid=2530087)
- Patel, T. R., Gould, G. C., Baehring, J. M., & Piepmeyer, J. M. (2012). Surgical approaches to lateral and third ventricular tumors. In *Schmidke and Sweet operative neurosurgical techniques: Indications, methods, and results: Sixth edition* (pp. 330–338). Elsevier, Inc. <https://doi.org/10.1016/b978-1-4160-6839-6.10027-9>
- 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. <https://doi.org/10.48550/arXiv.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/>



- Reuter, M., Rosas, H. D., & Fischl, B. (2010). Highly accurate inverse consistent registration: A robust approach. *NeuroImage*, 53, 1181–1196. <https://doi.org/10.1016/j.neuroimage.2010.07.020>
- Rodrigues, L., Rezende, T., Zanesco, A., Hernandez, A. L., Franca, M., & Rittner, L. (2020). Hypothalamus fully automatic segmentation from MR images using a U-Net based architecture. In *15th International Symposium on Medical Information Processing and Analysis* (p. 113300J). International Society for Optics and Photonics. <https://doi.org/10.1117/12.2542585>
- Rodrigues, L., Rezende, T. J. R., Wertheimer, G., Santos, Y., França, M., & Rittner, L. (2022). A benchmark for hypothalamus segmentation on T1-weighted MR images. *NeuroImage*, 264, 119741. <https://doi.org/10.1016/j.neuroimage.2022.119741>
- Roh, J. H., Jiang, H., Finn, M. B., Stewart, F. R., Mahan, T. E., Cirrito, J. R., Heda, A., Snider, B. J., Li, M., Yanagisawa, M., de Lecea, L., & Holtzman, D. M. (2014). Potential role of orexin and sleep modulation in the pathogenesis of Alzheimer's disease. *Journal of Experimental Medicine*, 211, 2487–2496. <https://doi.org/10.1084/jem.20141788>
- 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* (pp. 234–241). Springer. [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
- Roy, A. G., Conjeti, S., Navab, N., Wachinger, C., & Alzheimer's Disease Neuroimaging Initiative. (2019). QuickNat: A fully convolutional network for quick and accurate segmentation of neuroanatomy. *NeuroImage*, 186, 713–727. <https://doi.org/10.1016/j.neuroimage.2018.11.042>
- Rushmore, R. J., Sunderland, K., Carrington, H., Chen, J., Halle, M., Lasso, A., Papadimitriou, G., Prunier, N., Rizzoni, E., Vessey, B., Wilson-Braun, P., Rath, Y., Kubicki, M., Bouix, S., Yeterian, E., & Makris, N. (2022). Anatomically curated segmentation of human subcortical structures in high resolution magnetic resonance imaging: An open science approach. *Frontiers in Neuroanatomy*, 16. <https://doi.org/10.3389/fnana.2022.894606>
- Saper, C. B., & Lowell, B. B. (2014). The hypothalamus. *Current Biology*, 24, R1111–R1116. <https://doi.org/10.1016/j.cub.2014.10.023>
- Schindler, S., Schönknecht, P., Schmidt, L., Anwender, A., Strauß, M., Trampel, R., Bazin, P.-L., Möller, H. E., Hegerl, U., Turner, R., & Geyer, S. (2013). Development and evaluation of an algorithm for the computer-assisted segmentation of the human hypothalamus on 7-tesla magnetic resonance images. *PLoS One*, 8, e66394. <https://doi.org/10.1371/journal.pone.0066394>
- Shapiro, N. L., Todd, E. G., Billot, B., Cash, D. M., Iglesias, J. E., Warren, J. D., Rohrer, J. D., & Bocchetta, M. (2022). In vivo hypothalamic regional volumetry across the frontotemporal dementia spectrum. *NeuroImage: Clinical*, 35, 103084. <https://doi.org/10.1016/j.nicl.2022.103084>
- Sorensen, T. A. (1948). A method of establishing groups of equal amplitude in plant sociology based on similarity of species content and its application to analyses of the vegetation on Danish commons. *Biologiske Skrifter*, 5, 1–34. [https://www.royalacademy.dk/Publications/High/295\\_S%C3%B8rensen,%20Thorvald.pdf](https://www.royalacademy.dk/Publications/High/295_S%C3%B8rensen,%20Thorvald.pdf)
- 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)*. <https://cds.ismrm.org/protected/16MProceedings/PDFfiles/6865.html>
- Sudre, C. H., Cardoso, M. J., Ourselin, S., & for the Alzheimer's Disease Neuroimaging Initiative. (2017). Longitudinal segmentation of age-related white matter hyperintensities. *Medical Image Analysis*, 38, 50–64. <https://doi.org/10.1016/j.media.2017.02.007>
- Taha, A. A., & Hanbury, A. (2015). Metrics for evaluating 3D medical image segmentation: Analysis, selection, and tool. *BMC Medical Imaging*, 15, 1–28. <https://doi.org/10.1186/s12880-015-0068-x>
- Thomas, K., Beyer, F., Lewke, G., Zhang, R., Schindler, S., Schönknecht, P., Stumvoll, M., Villringer, A., & Witte, A. V. (2019). Higher body mass index is linked to altered hypothalamic microstructure. *Scientific Reports*, 9, 1–11. <https://doi.org/10.1038/s41598-019-53578-4>
- van der Kouwe, A. J., Benner, T., Salat, D. H., & Fischl, B. (2008). Brain morphometry with multiecho mprage. *NeuroImage*, 40, 559–569. <https://doi.org/10.1016/j.neuroimage.2007.12.025>
- Van Essen, D. C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T. E., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S. W., Della Penna, S., Feinberg, D., Glasser, M. F., Harel, N., Heath, A. C., Larson-Prior, L., Marcus, D., Michalareas, G., Moeller, S., ... Yacoub, E. (2012). The human connectome project: A data acquisition perspective. *NeuroImage*, 62, 2222–2231. <https://doi.org/10.1016/j.neuroimage.2012.02.018>
- Van Leemput, K., Maes, F., Vandermeulen, D., & Suetens, P. (1999). Automated model-based tissue classification of MR images of the brain. *IEEE Transactions on Medical Imaging*, 18, 897–908. <https://doi.org/10.1109/42.811270>
- Van Tulder, G., & de Bruijne, M. (2015). Why does synthesized data improve multi-sequence classification? In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part I* 18 (pp. 531–538). Springer. [https://doi.org/10.1007/978-3-319-24553-9\\_65](https://doi.org/10.1007/978-3-319-24553-9_65)
- van Wamelen, D. J., & Aziz, N. A. (2021). Hypothalamic pathology in Huntington disease. *Handbook of Clinical Neurology*, 182, 245–255. <https://doi.org/10.1016/b978-0-12-819973-2.00017-4>
- Varsavsky, T., Eaton-Rosen, Z., Sudre, C. H., Nachev, P., & Cardoso, M. J. (2018). Pimms: Permutation invariant multi-modal segmentation. In *Deep learning in medical image analysis and multimodal learning for clinical decision support* (pp. 201–209). Springer. [https://doi.org/10.1007/978-3-030-00889-5\\_23](https://doi.org/10.1007/978-3-030-00889-5_23)
- Wilcoxon, F. (1992). Individual comparisons by ranking methods. In S. Kotz & N. L. Johnson (Eds.), *Breakthroughs in statistics* (pp. 196–202). Springer. [https://doi.org/10.1007/978-1-4612-4380-9\\_16](https://doi.org/10.1007/978-1-4612-4380-9_16)
- Wolters, A., Heijmans, M., Michielse, S., Leentjens, A., Postma, A., Jansen, J., Ivanov, D., Duits, A., Temel, Y., & Kuijf, M. (2020). The TRACK-PD study: Protocol of a longitudinal ultra-high field imaging study in Parkinson's disease. *BMC Neurology*, 20, 1–10. <https://doi.org/10.1186/s12883-020-01874-2>



## APPENDIX A

**Appendix Table A1.** Sequence parameters for the T1-weighted and T2-weighted versions in the Rhineland Study.

T1w sequence			T2w sequence			
Parameters	Version		Parameters	Version		
	$T1w^a$	$T1w^b$		$T2w^a$	$T2w^b$	$T2w^c$ $T2w^d$
Repetition time (TR)	2560 ms		Repetition time (TR)	2800 ms		
Inversion time (TI)	1100 ms		Echo time (TE)	405 ms		
Flip angle	7�		Matrix size	320 � 320 � 224		
Matrix size	320 � 320 � 224		Phase-encoding direc.**	A>P	R>L	A>P A>P
PI acc. Factor	1�3	1�2	PI acc. factor	3�1	2�1	1�2 <sup>+</sup>
Readout bandwidth	240 Hz/pixel	740 Hz pixel	PI ref. scan	Integrated		External
Echo time (TE)	2.94 ms <sup>*</sup>	1.68 ms to 6.51 ms <sup>**</sup>	Acquisition time (TA)	3:57 minutes	4:30 minutes	4:47 minutes
Acquisition time (TA)	3:43 minutes	6:35 minutes				

To date, there have been two versions of the T1w sequence ( $T1w^{a-b}$ ) and four versions of the T2w sequence ( $T2w^{a-d}$ )—care was taken to preserve the image contrast between versions for both sequences.

\*1 echo, \*\*4 echoes combined to 1.

+with one CAIPIRINHA shift (Breuer et al., 2006), \*\* A: anterior, P: posterior, R: right, and L: Left.

**Appendix Table A2.** Demographics of the Rhineland Study participants for all different datasets.

	Case study	Test-retest	In-house	Total	p-value
	(N = 463)	(N = 21)	(N = 50)	(N = 534)	
Sex					0.801
Women	276 (59.6%)	11 (52.4%)	30 (60.0%)	317 (59.4%)	
Men	187 (40.4%)	6 (47.6%)	20 (40.0%)	217 (40.6%)	
Age					0.805
Mean (SD)	54.9 (14.2)	56.4 (9.3)	54.0 (15.2)	54.9 (14.1)	
Range	30.0 - 95.0	40.0 - 74.0	31.0 - 79.0	30.0 - 95.0	
T1w version					0.061
a	71 (15.3%)	0 (0.0%)	4 (8.0%)	75 (14.0%)	
b	392 (84.7%)	21 (100.0%)	46 (92.0%)	459 (86.0%)	
T2w version					< 0.001
a	71 (15.3%)	0 (0.0%)	4 (8.0%)	75 (14.0%)	
b	14 (3.0%)	0 (0.0%)	2 (4.0%)	16 (3.0%)	
c	269 (58.1%)	0 (0.0%)	27 (54.0%)	296 (55.4%)	
d	109 (23.5%)	21 (100.0%)	17 (34.0%)	147 (27.5%)	

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.

**Appendix Table A3.** Demographics for the training and testing in-house dataset.

	Trainset				Testset	Total	p-value
	Split_1 (N = 11)	Split_2 (N = 11)	Split_3 (N = 11)	Split_4 (N = 11)	(N = 6)	(N = 50)	
Sex							0.857
Women	6 (54.5%)	7 (63.6%)	8 (72.7%)	6 (54.5%)	3 (50.0%)	30 (60.0%)	
Men	5 (45.5%)	4 (36.4%)	3 (27.3%)	5 (45.5%)	3 (50.0%)	20 (40.0%)	
Age							0.439
Mean (SD)	46.7 (14.8)	53.5 (16.0)	56.5 (15.3)	58.5 (15.0)	55.2 (14.9)	54.0 (15.2)	
Range	31.0 - 69.0	31.0 - 77.0	32.0 - 79.0	35.0 - 76.0	35.0 - 71.0	31.0 - 79.0	

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.

**Appendix Table A4.** Demographics for the UK Biobank participants for all different datasets.

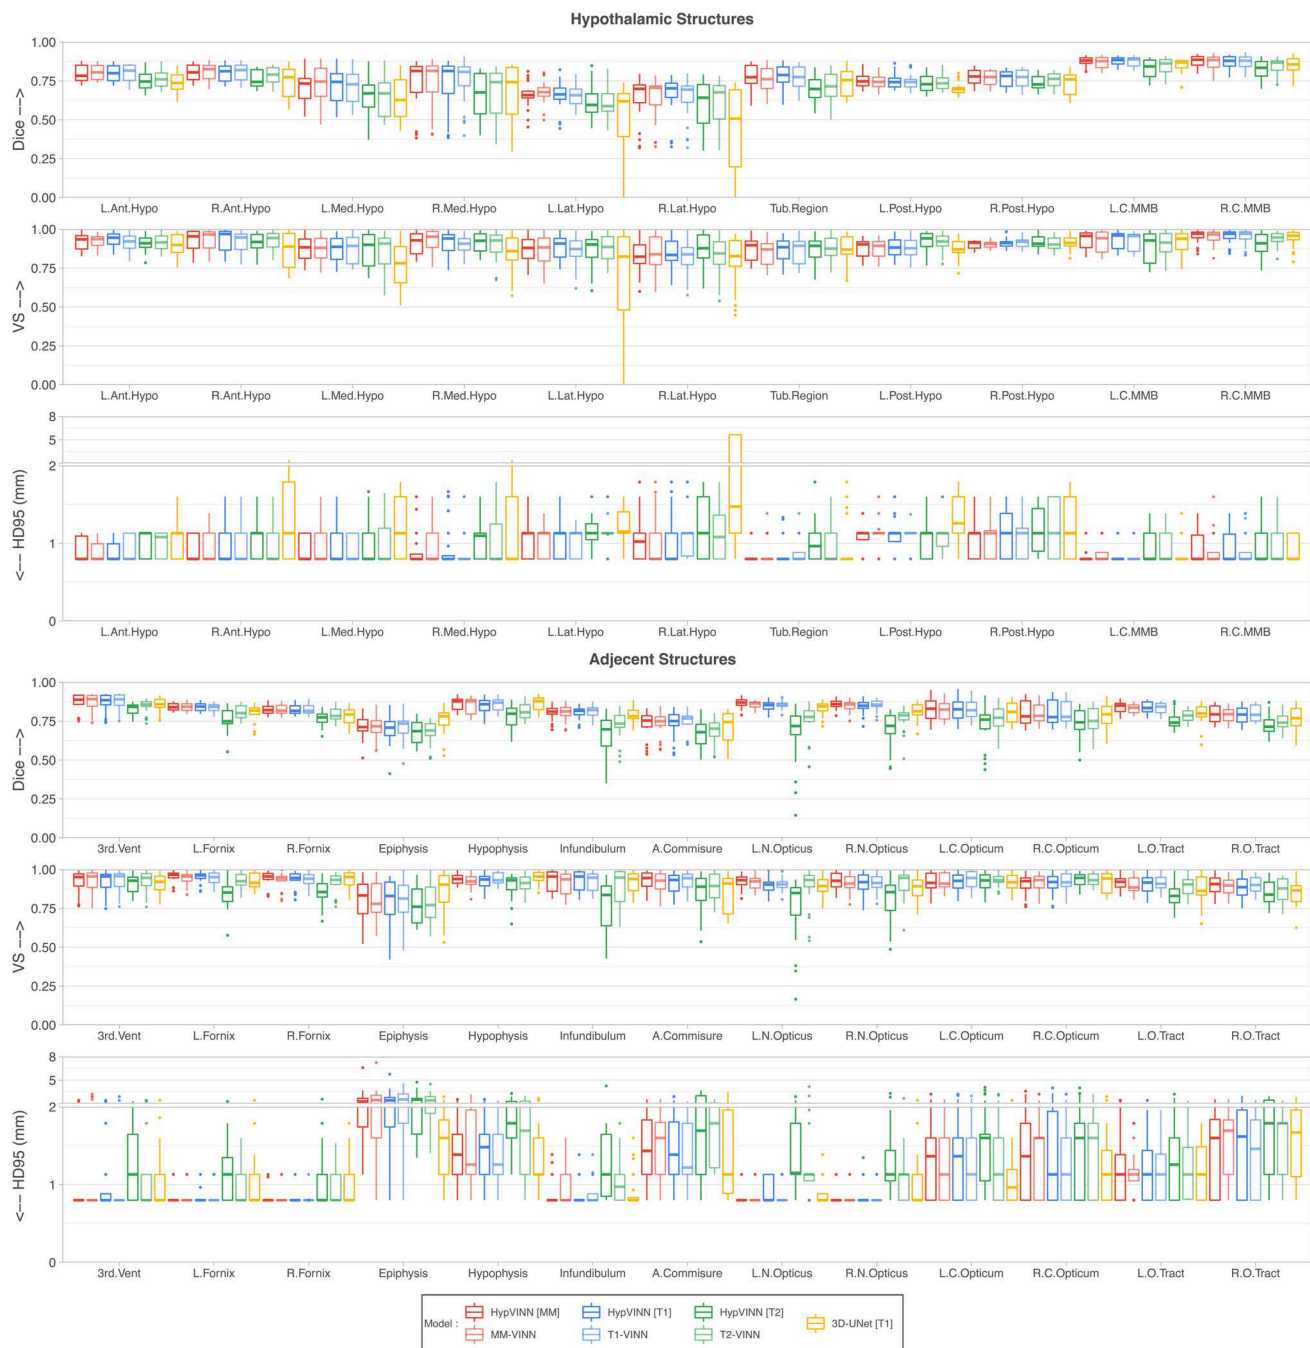
	Case Study	Generalizability	Total	p-value
	(N = 535)	(N = 9)	(N = 544)	
Sex				0.857
Women	281 (52.5%)	5 (55.6%)	286 (52.6%)	
Men	254 (47.5%)	4 (44.4%)	258 (47.4%)	
Age				0.050
Mean (SD)	63.9 (7.7)	58.7 (11.3)	63.8 (7.8)	
Range	46.0 - 82.0	45.0 - 77.0	45.0 - 82.0	

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.


**Appendix Table A5.** Test-retest reliability: Intra-class correlation (ICC) with a 95% confidence interval and volume similarity (VS) between volume estimates across sequences in a test-retest scenario for the 21 cases of the test-retest dataset.

Model	Hypothalamic			Others			Optic		
	ICC(A,1)		VS	ICC(A,1)		VS	ICC(A,1)		VS
	ICC [95% CI]	Mean (SD)		ICC [95% CI]	Mean (SD)		ICC [95% CI]	Mean (SD)	
Only T1w input									
a: T1-VINN	0.984 [0.959 - 0.994]	0.990 (0.011)		0.997 [0.993 - 0.999]	0.993 (0.006)		0.982 [0.953 - 0.993]	0.994 (0.005)	
b: HypVINN	0.982 [0.953 - 0.993]	0.987 (0.025)		0.999 [0.997 - 1.000]	0.996 (0.003)		0.985 [0.955 - 0.994]	0.994 (0.005)	
(Ours)									
Multi-modal (MM) input (T1w & T2w)									
c: MM-VINN	0.990 [0.975 - 0.996]	0.990 (0.010)		0.998 [0.995 - 0.999]	0.994 (0.006)		0.972 [0.879 - 0.990]	0.992 (0.006)	
d: HypVINN	0.984 [0.957 - 0.994]	0.989 (0.015)		0.999 [0.998 - 1.000]	0.996 (0.003)	a	0.986 [0.955 - 0.995]	0.994 (0.004)	
(Ours)									

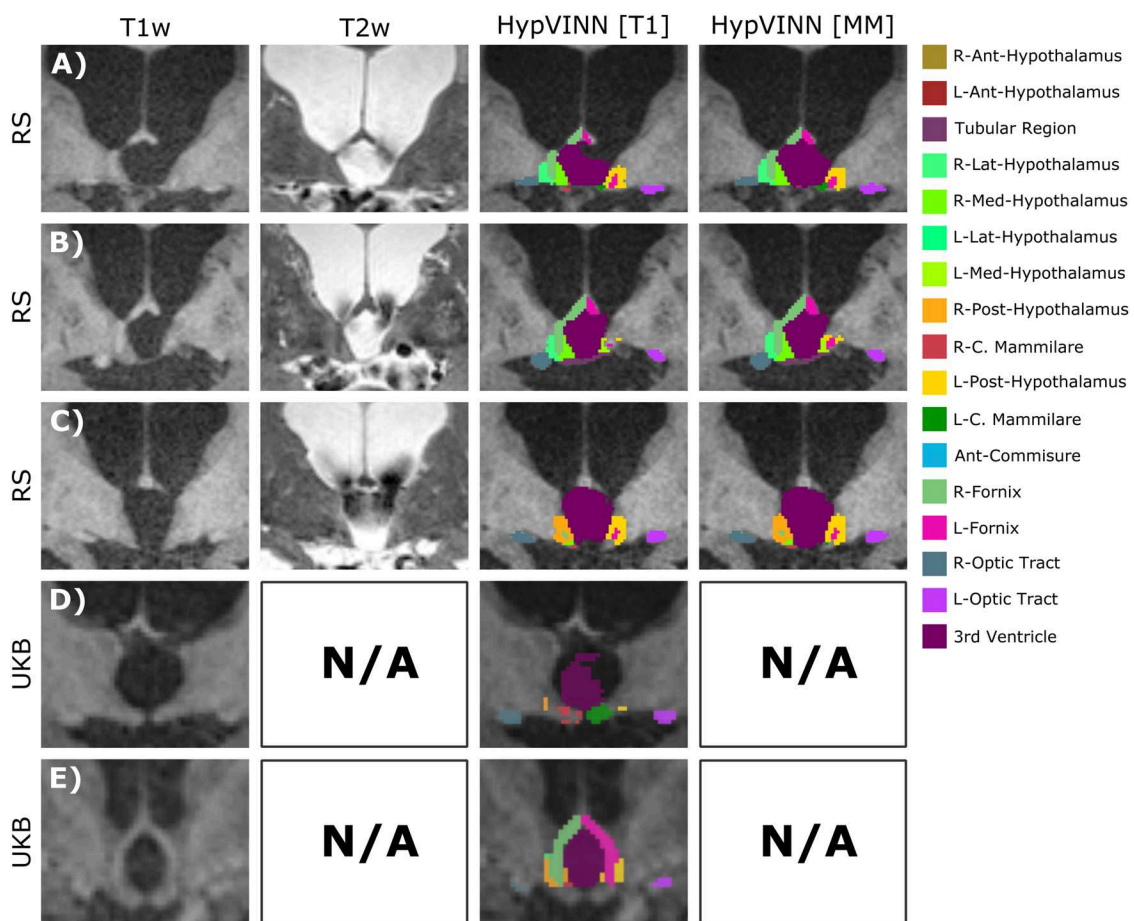
All automated methods exhibit excellent test-retest agreement between in-session volume estimates. Note: the statistical significance column (Signif.) indicates which other models the model outperforms (Wilcoxon signed-rank test, corrected  $p < 0.05$ ).



**Appendix Fig. A1.** Per structure segmentation performance of the F-CNN models on the unseen in-house test-set. We observe that models with a T1w image as part of its input have comparable results in all structures to the global ones. However, there is a slight decrease in Dice performance in the medial and lateral hypothalamus ( $\text{Dice} < 0.75$ ) compared to the other hypothalamic structures for the 2.5D models. For the 3D model, a similar trend is also observed in the medial hypothalamus; however, in the lateral hypothalamus, performance drastically diminishes in all evaluation metrics ( $\text{Dice} < 0.5$ ,  $\text{VS} < 0.8$ , and  $\text{HD95} > 1.2$  mm). Furthermore, for the adjacent hypothalamic structures, all 2.5D models present difficulties in localizing the epiphysis and recognizing its boundaries ( $\text{Dice} \leq 0.75$ ,  $\text{VS} \leq 0.8$ , and  $\text{HD95} \geq 2$  mm). Moreover, the epiphysis is the only structure from the 24 segmented ones where the 3D model outperforms the T1 and multi-modal comparative baselines ( $\text{Dice} = 0.7558$ ,  $\text{VS} = 0.8571$ , and  $\text{HD95} = 1.6386$  mm). Finally, using a T2w scan as the only source for inferring information is consistently underperforming in all structures, especially in the optic region (e.g., optic nerve) and middle hypothalamic region (e.g., medial and lateral hypothalamus and tubular region). Nonetheless, the inclusion of T2w into the current segmentation task appears to be beneficial as *HypVINN* with multi-modal input outperforms its T1w-only counterpart in most structures ( $\text{Dice}$ : 16/24,  $\text{VS}$ : 14/24, and  $\text{HD95}$ : 18/24).

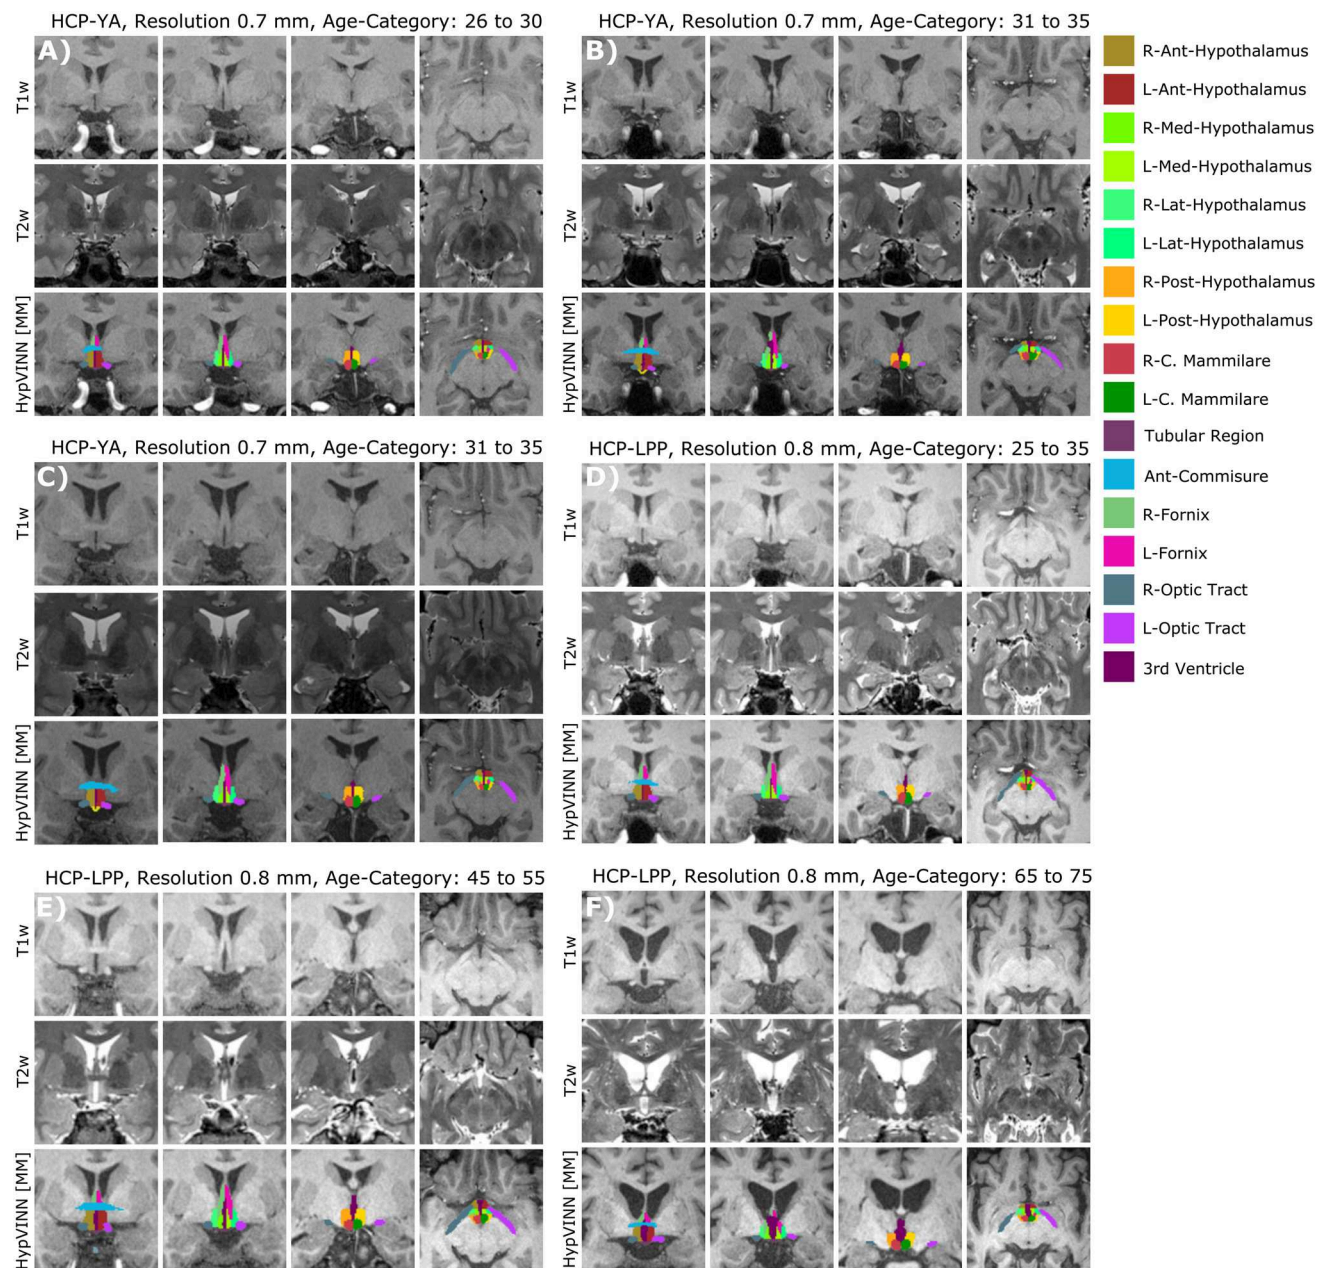


**Appendix Fig. A2.** T1-Block learnable modality weight during training. The T1-block has a much higher value ( $\approx 0.75$ ) than the T2-block weight ( $\approx 0.25$ ) in *HypVINN*'s fusion module, starting in the early training steps in all four cross-validation training splits (i.e., S1, S2, S3, and S4). Thus, performance is mainly driven by the T1-derived information, with T2w being only a support modality.



**Appendix Fig. A3.** Examples of excluded cases from the Rhineland Study (RS) and UK Biobank (UKB) after visual quality assessment. (A-E) Unclear boundary of the hypothalamus due to severe enlargements of the third ventricle (i.e., out-of-distribution cases) producing segmentation errors. Note: each row represents a different participant with corresponding MRI modalities (T1-weighted (T1w) and T2w-weighted (T2w)—if available), and automated generated segmentations on the coronal view. The color scheme for the visible structures is presented on the right.





**Appendix Fig. A4.** Examples of correct predictions in the Human Connectome Project (HCP) young adults (HCP-YA, A-C) and HCP lifespan pilot project (HCP-LPP, D-E) datasets (Bookheimer et al., 2019; Harms et al., 2018; Van Essen et al., 2012) from our proposed *HypVINN* with multi-modal input (MM) for six random participants. We observe that our tool shows promising results in both available HCP resolutions (0.7 mm and 0.8 mm). Furthermore, our tool seems to generalize well across age categories inside the training age range (training data started at age 30). However, all the above observations are only qualitative, and no accuracy segmentation metrics can be computed as manual annotations are unavailable for this dataset. Note: T1w, T2w, and *HypVINN* outcomes are presented for each participant. Furthermore, in each participant's row, the first three images display the different hypothalamic structures on the coronal view, and the remaining images show the structures on the axial view. The color lookup table for all visible structures is presented on the right.

## APPENDIX B. ABLATION ANALYSIS

We execute ablation analysis to optimize the fusion module weighting scheme inside the *HM-VINN* architecture by training the model with global and per-channel modality weights. First, all networks are trained from scratch using the four data-splits from the in-house training-set in a leave-one-out cross-validation approach. Afterwards, the best model is chosen based on the cross-validation performance in the hold-out validation-sets. The three evaluation metrics (Dice, VS, HD95) are computed per input modality combination (i.e., only T1w or only T2w, or

both) between the predicted maps after view aggregation and manual labels. Finally, improvements in segmentation performance are confirmed by statistical testing (corrected  $p < 0.05$ ).

We observe that utilizing global weights outperforms per-channel weights in all comparative metrics and all inference scenarios with statistical significance for the standalone T2w input in all three metrics and for the T1w & T2w input only in Dice, as presented in [Appendix Table B1](#). Therefore, we utilize the global weighting scheme as the fusion module configuration for this work.

**Appendix Table B1.** Fusion module weighting scheme optimization: Mean (and standard deviation) of segmentation performance metrics per input modality of the ablative *hetero-modal VINN* (*HM-VINN*) architectures on the validation set.

Only T1w input							
Experimental setup		Dice↑		VS↑		HD95 (mm)↓	
Model	Weighting Scheme	Mean (SD)	Signif.	Mean (SD)	Signif.	Mean (SD)	Signif.
a: HM-VINN	Global	<b>0.8068 (0.0841)</b>		<b>0.9164 (0.0748)</b>		<b>1.0916 (0.8579)</b>	
b: HM-VINN	Per Channel	0.8042 (0.0864)		0.9160 (0.0753)		1.0953 (0.7277)	
Only T2w input							
Experimental setup		Dice↑		VS↑		HD95 (mm)↓	
Model	Weighting scheme	Mean (SD)	Signif.	Mean (SD)	Signif.	Mean (SD)	Signif.
a: HM-VINN	Global	<b>0.7354 (0.1115)</b>	<sup>b</sup>	<b>0.8753 (0.1166)</b>	<sup>b</sup>	<b>1.4154 (1.3291)</b>	<sup>b</sup>
b: HM-VINN	Per Channel	0.7119 (0.1236)		0.8424 (0.1424)		1.700 (2.3105)	
T1w & T2w input							
Experimental Setup		Dice↑		VS↑		HD95 (mm)↓	
Model	Weighting scheme	Mean (SD)	Signif.	Mean (SD)	Signif.	Mean (SD)	Signif.
a: HM-VINN	Global	<b>0.8128 (0.0814)</b>	<sup>b</sup>	<b>0.9202 (0.0706)</b>		<b>1.0508 (0.6965)</b>	
b: HM-VINN	Per Channel	0.8079 (0.0869)		0.9187 (0.0754)		1.0678 (0.7445)	

Global weights outperform per-channel weights in all comparative metrics and all inference scenarios. Note: the statistical significance column (Signif.) indicates which other models the model outperforms (Wilcoxon signed rank test, corrected  $p < 0.05$ ), and bold values represent the best model per input modality combination.

## APPENDIX C. CRITERIA FOR MANUAL ANNOTATION OF HYPOTHALAMIC ADJACENT STRUCTURES

In Appendix Tables C1 and C2, we present the criteria for manual annotation of hypothalamic adjacent structures and sub-regions in T1w and T2w images. The sup-

port of a T2w image was omitted for segmenting UK Biobank data as these data were unavailable. Furthermore, no protocol modification was carried out due to the differences in data resolution—Rhineland Study 0.8 mm isotropic resolution and UK Biobank 1 mm isotropic resolution.

**Appendix Table C1.** Criteria for manual annotation of hypothalamic adjacent structures.

Structure	Bilateral <sup>*</sup>	Labeling <sup>**</sup>	Note
Optic system	Yes	The optic system is composed of the optic nerves, tracts, and chiasms. The optic chiasm was separated from the optic nerves and tracts at an angle orthogonal to the chiasm at the optic nerve–chiasm and optic tract–chiasm junctions, respectively (Avery et al., 2016).	Using axial T1-weighted images.
Anterior commissure	No	A thick fiber bundle above the 3 <sup>rd</sup> ventricle and underneath the anterior horns of the lateral ventricles. It can easily be identified using the brain ventricles and optic tracts as landmarks (Güngör et al., 2017).	Labeling on coronal sections in the rostro-caudal direction on T1-weighted images.
Fornices	Yes	Thick white matter fiber bundles that were labeled in the area where they touch the anterior commissure rostrally and merge with the mammillary bodies caudally; this part of the fornix is generally referred to as the “columna fornicis.”	Using coronal sections of T1-weighted sequences.
Hypophysis (i.e., the pituitary gland)	No	A relatively round structure inferior to the 3 <sup>rd</sup> ventricle and rostral to the brain stem, occupying the sella turcica.	Using sagittal, axial, and coronal sections of T1- and T2-weighted images.
Infundibulum (i.e., the pituitary stalk)	No	The stalk-like structure that connects the hypophysis to the hypothalamus.	Labeling was done on coronal sections by moving caudally from the posterior commissure, with its contours demarcated by its pine-like shape and the surrounding cerebrospinal fluid.
Epiphysis (i.e., the pineal gland)	No	A low-intensity (on T1-weighted images), pine-shaped unpaired midline brain structure that lies between the caudal recess of the third ventricle and the quadrigeminal cistern (Park et al., 2020).	
3 <sup>rd</sup> ventricle	No	Anterior border: lamina terminalis.  Lateral border: hypothalamus and thalamus. Superior border: the roof of the third ventricle starts anteriorly at the foramen of Monro and ends posteriorly in the suprapineal recess. Posterior border: the posterior commissure, the pineal body, the habenular commissure, and the suprapineal recess above (Patel et al., 2012). Inferior border: formed from anterior to posterior by the optic recess, the infundibular recess, the tuber cinereum, the mammillary bodies, and the posterior perforated substance (Chaichana & Quinones-Hinojosa, 2019).	Using sagittal, axial, and coronal sections of T1- and T2-weighted images.

\*Bilateral structures were defined as those regions that could be separated into a (non-contiguous) left and right half with respect to the mid-sagittal plane.

\*\*Labeling was mainly done using T1-weighted images, unless specified otherwise.

**Appendix Table C2.** Criteria for manual annotation of hypothalamic sub-regions.

Structure	Bilateral*	Labeling**	Note
Anterior hypothalamus	Yes	Medial border: 3 <sup>rd</sup> ventricle. Lateral border: lateral border of the optic tract and the other adjacent white matter tracts (Lemaire et al., 2011). Anterior border: lamina terminalis attached to the optic chiasm. Posterior border: vanishment of the anterior commissure on coronal sections in the rostro-caudal direction (coinciding with the coronal plane through the posterior border of the anterior commissure and the anterior tip of the infundibulum). Superior border: horizontal plane through the anterior commissure. Inferior border: optic chiasm and infundibulum (Dudás, 2021).	The <b>supraoptic nuclei</b> were included in this region and were not labeled separately as the spatial resolution was too low for accurate segmentation of these small structures.
Medial hypothalamus	Yes	Medial border: 3 <sup>rd</sup> ventricle. Lateral border: fornices. Anterior border: vanishment of the anterior commissure on coronal sections in the rostro-caudal direction. Posterior border: appearance of the mammillary bodies on coronal sections in the rostro-caudal direction. Superior border: the diencephalic fissure. Inferior border: the boundaries of the tuberal region underneath (Makris et al., 2013).	
Lateral hypothalamus	Yes	Medial border: fornices. Lateral border: optic tract and the other adjacent white matter tracts. Anterior border: vanishment of the anterior commissure on coronal sections in the rostro-caudal direction. Posterior border: appearance of the mammillary bodies on coronal sections in the rostro-caudal direction. Superior border: the diencephalic fissure. Inferior border: the boundaries of the tuberal region and basal cistern underneath.	
Posterior hypothalamus	Yes	Medial border: 3 <sup>rd</sup> ventricle. Lateral border: white matter tracts. Anterior border: appearance of the mammillary bodies on coronal sections in the rostro-caudal direction. Posterior border: vanishment of the mammillary bodies on coronal sections in the rostro-caudal direction. Superior border: horizontal plane through the diencephalic fissure. Inferior border: boundaries with the mammillary bodies below.	
Tubular region	No	The area was defined as the region underneath the 3 <sup>rd</sup> ventricle and enclosed by the mammillary bodies caudally and the anterior hypothalamus rostrally, with its superior and inferior borders on each side defined by the horizontal planes going through the superior border of the floor of the third ventricle and the interpeduncular cistern, respectively. Median eminence: the protuberant region between the unpaired infundibular nucleus and the mammillary bodies that had a low intensity on the sagittal view on T2-weighted sequences. Infundibular nucleus: dorsocaudal to the junction of the infundibulum (i.e., the pituitary stalk) and the hypothalamus, and was labeled on the sagittal view using T1 (high-intensity) and T2 (low-intensity) weighted images. Tubero-mammillary nucleus: the remaining areas in the tuberal region.	
Mammillary bodies	Yes	Two small, rounded structures at the caudal end of the 3 <sup>rd</sup> ventricle. These structures were labeled using both coronal sections in the rostro-caudal direction and axial sections in the dorso-medial direction on T1-weighted images.	

\*Bilateral structures were defined as those regions that could be separated into a (non-contiguous) left and right half with respect to the mid-sagittal plane.

\*\*Labeling was mainly done using T1-weighted images, unless specified otherwise.