



FastSurferVINN: Building resolution-independence into deep learning segmentation methods—A solution for HighRes brain MRI

Leonie Henschel^{a,1}, David K  gler^{a,1}, Martin Reuter^{a,b,c,*}

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

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

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

ARTICLE INFO

Keywords:

Computational neuroimaging
Deep learning
Structural MRI
Artificial intelligence
High-resolution

ABSTRACT

Leading neuroimaging studies have pushed 3T MRI acquisition resolutions below 1.0 mm for improved structure definition and morphometry. Yet, only few, time-intensive automated image analysis pipelines have been validated for high-resolution (HiRes) settings. Efficient deep learning approaches, on the other hand, rarely support more than one fixed resolution (usually 1.0 mm). Furthermore, the lack of a standard submillimeter resolution as well as limited availability of diverse HiRes data with sufficient coverage of scanner, age, diseases, or genetic variance poses additional, unsolved challenges for training HiRes networks. Incorporating resolution-independence into deep learning-based segmentation, i.e., the ability to segment images at their native resolution across a range of different voxel sizes, promises to overcome these challenges, yet no such approach currently exists. We now fill this gap by introducing a Voxel-size Independent Neural Network (VINN) for resolution-independent segmentation tasks and present FastSurferVINN, which (i) establishes and implements resolution-independence for deep learning as the first method simultaneously supporting 0.7–1.0 mm whole brain segmentation, (ii) significantly outperforms state-of-the-art methods across resolutions, and (iii) mitigates the data imbalance problem present in HiRes datasets. Overall, internal resolution-independence mutually benefits both HiRes and 1.0 mm MRI segmentation. With our rigorously validated FastSurferVINN we distribute a rapid tool for morphometric neuroimage analysis. The VINN architecture, furthermore, represents an efficient resolution-independent segmentation method for wider application.

1. Introduction

While neuroimaging pipelines have benefited substantially from the standardization of Magnetic Resonance Imaging (MRI) at 1.0 mm, the resulting fixed-resolution paradigm now hinders transition to high-resolution (HiRes) MRI. With the hope of advancing quantification of structural detail, increasing explanatory power, and improving our understanding of the brain in health and disease (Glasser et al., 2013; Mellerio et al., 2014; Solano-Castiella et al., 2011; Stankiewicz et al., 2011; Wattjes et al., 2006; Zaretskaya et al., 2018), leading large-cohort neuroimaging studies have started to acquire structural MRI at 3T field strength and 0.7–0.9 mm resolutions (see Section 2.1). However, the lack of reference segmentations and limited diversity of HiRes MRI datasets (e.g. regarding scanner, disease, genetic variation) lead to substantial limitations for bias-free method development and val-

idation. Additionally, since no de-facto standard resolution exists for HiRes imaging, neuroimaging tools introducing HiRes processing (Bazin et al., 2014; Gaser and Dahnke, 2016; Huntenburg et al., 2018; Yushkevich et al., 2014; Zaretskaya et al., 2018) have to provide resolution-independence instead of following the fixed-resolution paradigm. Although Convolutional Neural Networks (CNNs) deliver convincing performance under the fixed-resolution paradigm (Chen et al., 2018; Coup   et al., 2020; Henschel et al., 2020; Huo et al., 2019; Ito et al., 2019; McClure et al., 2019; Mehta et al., 2017; Roy et al., 2019; Sun et al., 2019; Wachinger et al., 2018), no methodological solution leverages explicit knowledge of the native image resolution, consequently limiting all output segmentations to one pre-defined voxel size and potentially ignoring important structural detail specifically for submillimeter scans. By introducing Voxel-size Independent Neural Networks (Voxel-size Independent Neural Networks (VINNs)), we now leverage the diversity

Abbreviations: HiRes, high-resolution; LowRes, low-resolution; VINN, Voxel-size Independent Neural Network; exSA, external scale augmentation; inSA, internal scale augmentation; SF, scale factor; PVE, partial volume effect.

* Corresponding author at: German Center for Neurodegenerative Diseases (DZNE), Bonn, Germany

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

¹ Equal Contribution.

<https://doi.org/10.1016/j.neuroimage.2022.118933>.

Received 8 October 2021; Received in revised form 22 December 2021; Accepted 24 January 2022

Available online 3 February 2022.

1053-8119/   2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>)

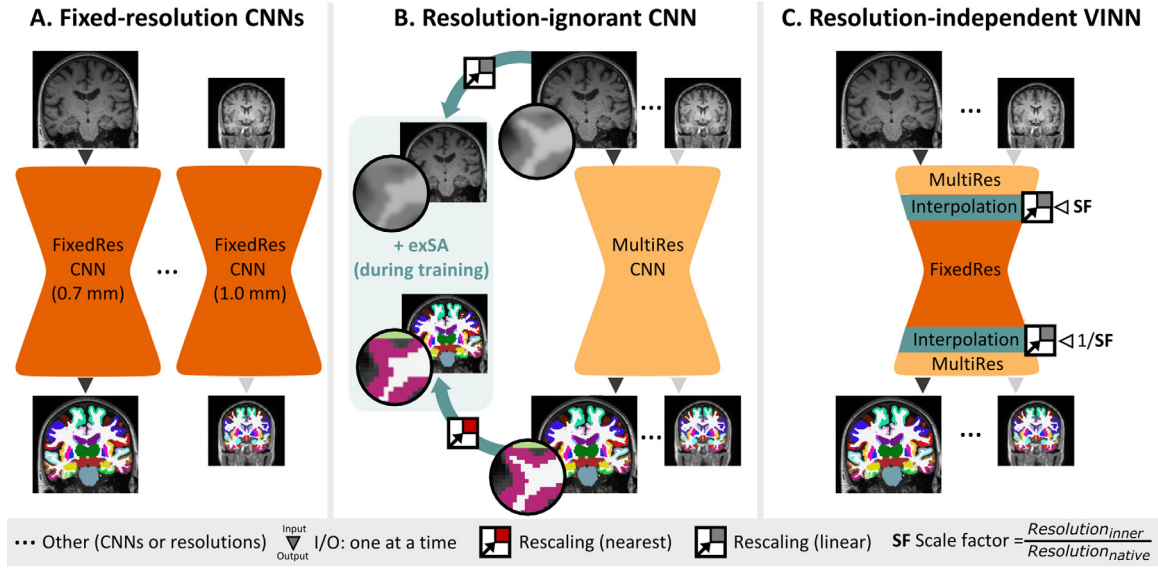


Fig. 1. Resolution-independence in deep learning networks: **A.** Dedicated fixed-resolution convolutional neural networks (CNNs) only work on the resolution they are trained on and are limited by the availability and quality of corresponding datasets. **B.** One single resolution-ignorant CNN can learn to segment multiple resolutions by training on a diverse dataset. External scale augmentation (+exSA, B. left) simulates resolutions with few or no training cases by resampling the image and the reference segmentation map. Here, however, lossy interpolation and resulting artefacts, especially from nearest-neighbour interpolation of discrete label maps, may result in a loss of structural details and sub-optimal performance. **C.** Our voxel size independent neural network (VINN) avoids interpolation of the images and discrete labels by integrating the interpolation step into the network architecture. Further, the explicit transition from the native resolution to a normalized internal resolution facilitates an understanding of the difference between image features (MultiRes blocks with distances measured in voxels) and anatomical features (FixedRes inner blocks with normalized distances).

of widely available 1.0 mm MRIs and enrich the model with details derived from HiRes MRI, achieving not only resolution-independence but improving segmentation performance across resolutions.

For resolution-independent deep learning, we establish two core requirements: 1. Native-resolution segmentation: the network's input and output for training and *crucially inference* should be at the native resolution, to avoid any external resampling. 2. Resolution-independence: the network should be able to learn and predict from images at a range of different resolutions. Segmentation for HiRes images additionally aims to improve the quality of fine structures (e.g. narrow sulci, gyri and white matter (WM) details). To achieve high-quality whole brain segmentation and avoid training biases, a neural network should generalize to various resolutions (seen and unseen during training) and datasets with different characteristics (e.g. with respect to scanners, demographics, diseases, genetic variation) ideally sharing and transferring knowledge between resolutions.

While some traditional neuroimaging pipelines fulfill these requirements (Bazin et al., 2014; Gaser and Dahnke, 2016; Huntenburg et al., 2018; Yushkevich et al., 2014; Zaretskaya et al., 2018), no related work directly addresses these challenges with deep learning. Fig. 1 illustrates two adaptations to form baseline solutions: A. dedicated fixed-resolution networks and B. a single resolution-ignorant CNN that accepts multiple resolutions through training. Training one dedicated network per resolution (A.) trades the potential of a larger, more diverse training corpus for compatibility with the native resolution raising bias and generalization limitations (there are, for example, currently no HiRes neurodegeneration datasets). However, a fixed-resolution network trained and evaluated on its native resolution represents an upper bound for achievable performance under same-size training datasets. On the other hand, if one network is naïvely trained on multiple resolutions (B.), each convolutional layer has to learn to generalize across scales as the network cannot easily differentiate between different voxel sizes (i.e. resolution-ignorant), which allocates network capacity to this task. To support this process and reduce potential bias from missing or unbalanced resolu-

tion data, one can add external scale augmentation (+exSA) by resampling individual images and reference labels during training. This, however, induces information loss and interpolation artefacts, e.g. from lossy nearest-neighbour (NN) interpolation of discrete label maps. Since neither approach (A. or B.) has been implemented or compared for sub-millimeter whole brain segmentation, we introduce respective baseline models utilizing our proposed, optimized micro-architecture in all models for a fair comparison.

To overcome limitations of both approaches (A. and B.), we propose a VINN with innovations including its micro-architecture and the addition of a HiRes loss. The core contribution, however, is the network-integrated resolution-normalization to support native segmentation at various voxel sizes. In fact, any UNet-based architecture (Ronneberger et al., 2015) is, by design, a multi-scale approach where pooling operations represent fixed-factor integer down- and up-scale transitions (usually by the scale factor 2). Our network-integrated resolution-normalization in VINN now replaces this fixed scale transition with a flexible re-scaling for the first and last scale transitions. This has the advantage of placing our interpolation operation at a position where information loss naturally occurs (down- or up-scaling via pooling). As illustrated in Fig. 1C., we retain compatibility with a range of resolutions for input and output (MultiRes) by shifting the interpolation into the architecture itself. At the same time, we leverage the lower variance of perceived size differences in the inner normalized resolution blocks (FixedRes). This has the advantages of (i) retaining important image information at the native resolution in the MultiRes blocks, (ii) interpolating multi-dimensional continuous feature maps rather than discrete labels or single slice images hence avoiding lossy NN interpolation and extending contextual neighbourhood information during the resolution-normalization, and (iii) disentangling perceived voxel versus actual structure size differences inside the VINN.

Especially the last point may have been underappreciated so far: Due to the nature of convolutional layers, CNNs 'perceive' distances and thus structure sizes by number of voxels rather than millimeters. Therefore,

the original voxel size impacts perceived distances for baseline architectures, requiring them to cover larger size variety during training and inference in the multi-resolution context. The VINN, on the other hand, transitions from a voxel-based distance context in the MultiRes blocks to a standardized distance context in FixedRes via the network-integrated resolution-normalization and, therefore, releases network capacity in the inner layers for other tasks.

1.1. Contributions

All in all, our VINN, for the first time, effectively addresses the challenges associated with HiRes MRI (i.e. multiple resolutions and reduced variety in datasets) in a single framework. We extensively test variations of the network architectures and demonstrate that inherent biases introduced by the unavailability of certain HiRes scans are reduced by transferring information across resolutions via our network-integrated resolution-normalization.

Specifically, we show that our VINN

- i. segments 3T brain MRI at their native resolutions (0.7–1.0 mm) into 95 regions in less than 1 min on the GPU and generalizes robustly to unseen resolutions within and beyond the training corpus,
- ii. significantly outperforms state-of-the-art scale augmentations as well as fixed-resolution models with respect to segmentation accuracy and generalizability with an optimized architecture, and
- iii. improves accuracy by combining and leveraging both the increased structural information from submillimeter 3T brain scans and the generalizability to intensity, scanner, disease, and other variations from standard 1.0 mm MRIs.

2. Related work

2.1. High-resolution MRI

In neuroimaging, spatial resolution is of great importance as the available voxel size directly dictates the degree to which fine-scale subcortical and cortical structures, specifically narrow gyri and sulci, can be resolved in an MRI. In particular, the diversity of tissue types in a single voxel, which heavily depends on its size, influences the signal intensity known as the partial volume effect (PVE). A strong partial volume effect (PVE) specifically complicates the delineation of tissue borders on the voxel grid (Glasser et al., 2013; Zaretskaya et al., 2018). HiRes images offer finer sampling of the underlying information, thereby directly reduce PVE and enable more accurate segmentations, improved volume-based measurements, morphometry, surface placement, and derived thickness measures (see Fig. 2) (Glasser et al., 2013; Luesebrink et al., 2013; Zaretskaya et al., 2018).

It is, therefore, not surprising that HiRes MRIs are becoming increasingly popular within the neuroimaging community. While most established large-scale neuroimaging studies acquire data at 1.0 mm, the current and next generation studies are shifting to submillimeter resolutions (e.g. Human Connectome Project (HCP) (Glasser et al., 2013), Rhineland Study (RS) (Breteler et al., 2014), Autism Brain Imaging Data Exchange II (ABIDE-II) (Di Martino et al., 2017), TRACK-PD (Wolters et al., 2020)). However, to date the consolidated superset of publicly available 3T HiRes neuroimages is sparse, unbalanced, and heterogeneous with respect to available submillimeter resolutions (e.g. HCP: 0.7 and 0.8, ABIDE-II: 0.7–0.9, RS: 0.8). Especially the limited data variety at each specific resolution poses a real challenge for data-driven computational methods, which translates into limited compatibility with scanners, age spans, and especially disease groups. Moreover, unbalanced training data can easily lead to the introduction of biases into the model.

While suffering from stronger PVEs and ensuing detrimental effects on segmentation accuracy, the available collection of standardized 1.0 mm images, on the other hand, is large and diverse. Consequently, a wide coverage of age-groups, diseases, genetic variants, and scanners




Fig. 2. Image resolution affects the detail of discrete segmentation label maps and derived measures such as surface models and thickness. **A.** The low-resolution image is less detailed and causes partial volume effects (PVEs) by accumulating signals across tissue boundaries into larger voxels, whereas **B.** High-resolution images and derived segmentations allow more precise region delineation and capture details, e.g. for improved shape or thickness analysis (white arrows).

can be retrieved from the rich reservoir of openly available MRI data sources (e.g. OpenfMRI database (Poldrack et al., 2013; Poldrack and Gorgolewski, 2014), OpenNeuro (Markiewicz et al., 2021), NITRC-IR²). Furthermore, manual reference labels for validation are exclusively published openly for 1.0 mm (Klein and Tourville, 2012). To address dataset sparsity and bias at submillimeter resolutions, we believe data-driven computational methods require built-in resolution-independence. Only a single model that spans across the available resolutions can simultaneously provide benefits for both HiRes and standard 1.0 mm image analysis.

2.2. Automated analysis of high-resolution images

Traditionally, common neuroimaging pipelines have been developed and optimized for 1.0 mm voxels (Fischl et al., 2002; Friston et al., 2007; Gaser and Dahnke, 2016; Jenkinson et al., 2012; Zhang et al., 2001) representing the de-facto standard for years. FreeSurfer offers a validated HiRes stream (Zaretskaya et al., 2018), which provides subsegmentation of the cortex into 31 structures per hemisphere (DKAtlas) (Klein and Tourville, 2012). However, one problem here is the extended processing time caused by the cubic voxel increase – a common issue limiting applicability of traditional tools to large cohort studies.

The introduction of CNNs for whole brain segmentation has substantially reduced processing times to seconds on the GPU. Recent works employ both 2.5D and 3D UNet architectures (Billot et al., 2020; Chen et al., 2018; Coupé et al., 2020; Henschel et al., 2020; Huo et al., 2019; Iglesias et al., 2021; Ito et al., 2019; McClure et al., 2019; Mehta et al., 2017; Roy et al., 2019; Sun et al., 2019; Wachinger et al., 2018). Since GPU memory limitations render full volume 3D models impractical specifically for higher number of feature channels and output classes, top performing methods process the volume in slices (QuickNat, FastSurfer (Henschel et al., 2020; Roy et al., 2019)) or in large patches (DeepNat, SLANT,

² www.nitrc.org.

AssemblyNet (Coupé et al., 2020; Huo et al., 2019; Wachinger et al., 2018)) and then leverage aggregation schemes to recombine predictions into the full volume.

Despite their success for brain segmentation and other applications no work has introduced deep learning for submillimeter whole brain segmentation. Since CNNs require a large and diverse number of volume and segmentation pairs for effective training, the missing availability of diverse training data hinders this resolution transition. Additionally, the cubic relationship between resolution increase and GPU memory requirements puts memory hungry 3D architectures at a disadvantage, e.g. factor of $2.92 = (1/0.7)^3$ memory demand for a 1 to 0.7 mm reduction of voxel sizes.

In the past, multi-branch segmentation frameworks (Fu et al., 2018; Gerard et al., 2020; Gu et al., 2018; Kamnitsas et al., 2017; Liu et al., 2020; van Rijthoven et al., 2021; Wang et al., 2017; Xu et al., 2021; Yang and Peng, 2018; Zheng et al., 2021) have been used to avoid memory issues while simultaneously leveraging HiRes information. Here, multiple potentially cross-linked pathways are dedicated to specific down-scaled or cropped versions of the original image. The same principle is employed in other scale-aware networks (Chen et al., 2016; Huang et al., 2021; Li et al., 2019) by implementing more trans-scale connections or combinations of different dilation and kernel sizes. To further avoid sub-optimal equal weighting of different scale information, the networks often include attention mechanisms (Chen et al., 2016; Qin et al., 2018; Xu et al., 2021; Yang and Peng, 2018; Zheng et al., 2021). In practice, these methods are only compatible with the discrete input and sub-level resolutions they were trained on, leaving subvoxel scaling more common in neuroimaging data (0.7/0.8/1) unexploited. In fact, their intent is to explicitly integrate information from multiple image scales rather than achieve multi-resolution compatibility. Critically, in neuroimaging settings the true scale is known (image resolution) while scale-aware architectures assume this knowledge is not available. In VINN we inject this explicit knowledge into the network at the spatial resolution normalization step.

2.3. Resolution-independence in deep learning

Built-in resolution-independence in CNNs has not been described for brain segmentation nor – to our knowledge – for any other segmentation tasks. Approaches such as (Billot et al., 2020; Iglesias et al., 2021) pre-sample input images (with associated reliability maps) to a common resolution (here 1.0 mm) and provide outputs there, which makes them inherently fixed-resolution techniques. While they can provide 1.0 mm segmentations (and even images) for lower resolutional clinical scans via heavy augmentation, they can neither profit from submillimeter details, nor provide native HiRes segmentations. A transfer to higher resolutions would require retraining with a fixed submillimeter training set. This is, however, problematic as (i) no standard submillimeter resolution exists, meaning one would either need to train multiple versions or focus on the highest available resolution to retain input resolution-independence, (ii) HiRes datasets, as mentioned before, demonstrate low subject variance (disease, age, genetics) making this approach susceptible to training biases, and (iii) upscaling the 3D-UNet architecture to 0.7 mm would require 2.9-times as much GPU memory surpassing memory limits.

Different from segmentation networks, state-of-the-art super-resolution networks aim specifically at reconstructing a (fixed resolution) HiRes image based on a low-resolution (LowRes) input, also requiring ground truth at the higher resolution. They often rely on pre-sampling, i.e. using interpolation-based up-sampling methods to initialize the desired output grid and fine-tune the features in the following network (Dong et al., 2016; Kim et al., 2016a; 2016b). Using augmentation, these networks are trained to restore images coming from various lower resolutions via a single model, similar to the discussed segmentation strategy (Fig. 3B.). While showing better performance than dedicated single-scale models (fixed input and fixed output resolution),

pre-sampling significantly increases the computational complexity and cost of the architectures due to the increase in image size. To overcome this limitation, interpolation-based post-sampling has been introduced recently (Alao et al., 2021; Shen et al., 2021). Here, the up-sampling step is shifted towards the end of the network architecture and performs the interpolation in the latent space. Interestingly, this post-sampling approach is as effective or even superior to the pre-sampling methods, while maintaining versatility with respect to the chosen output scale (Alao et al., 2021; Shen et al., 2021). All these super-resolution architectures, in addition to not being aimed at segmentation, differ from our approach by transitioning once, from the input resolution to an output resolution. In contrast, we insert two latent-space interpolation blocks transitioning both ways between the native and the inner resolution.

More generally, multi-source domain adaptations (MSDAs) leverage available data with different underlying distributions (i.e. resolutions in our case). This field of research focuses on enhancing the generalization ability of a model by transferring knowledge between resource-full source domains (i.e. LowRes data) to a sparsely represented target domain (i.e. HiRes data) by transforming either the features in the latent space or images on a pixel-level. In contrast to our resolution-normalization, the latent space transformations in MSDAs are based on optimizing a discrepancy (Guo et al., 2018; Peng et al., 2019; Zhu et al., 2019) or adversarial loss (Li et al., 2019; Wang et al., 2019; Xu et al., 2018). These latent space alignments are, however, often insufficient for segmentation tasks due to their focus on high-level information only (Zhao et al., 2019). To circumvent this problem, intermediate domain generators have been proposed and successfully applied for semantic segmentation (Hoffman et al., 2018; Russo et al., 2019; Zhao et al., 2019). Here, a pixel-level alignment between source and target domain is learned through Generative adversarial networks. However, these MSDA methods are limited to a single target distribution. Extension to multi-targets is a relatively unexplored area with only a few published methods for classification tasks so far (Chen et al., 2019; Gholami et al., 2020; Jin et al., 2020; Liu et al., 2020; Peng et al., 2019; Roy et al., 2021; Yang et al., 2020).

Spatial transformers (Jaderberg et al., 2015) represent another technology relying on internal interpolations as an important building block. Here, spatial invariance (registration) is targeted via a learnable affine transformation inside the network (localisation network). After the source feature map coordinates are computed (grid calculator), intensity values for each target pixel are determined via bi-linear interpolation (sampler). Spatial transformers hence attempt to implicitly learn data representations in a resolution-ignorant way. While our approach shares grid calculation and interpolation within the network with spatial transformers, our training approach is different: Instead of a localisation network, we directly determine the sampling-grid based on the input scale factors, i.e. the ratio between the native input resolution and the desired normalized inner resolution. Hence, we explicitly integrate knowledge about the image resolutions into the architecture. As a result, computational complexity is reduced while still achieving desired resolution-independence.

3. Material and methods

3.1. Datasets

The following three submillimeter MRI datasets were selected for training, testing, and validation of FastSurferVINN. An extended list of the used 1.0 mm datasets can be found in the Section A.3 including a tabulated overview of all used datasets (Table 3). When not specifically mentioned otherwise, all sets are balanced for gender, age, and study. Participants of the individual studies gave informed consent in accordance with the Institutional Review Board at each of the participating sites. Complete ethic statements are available at the respective study webpages.

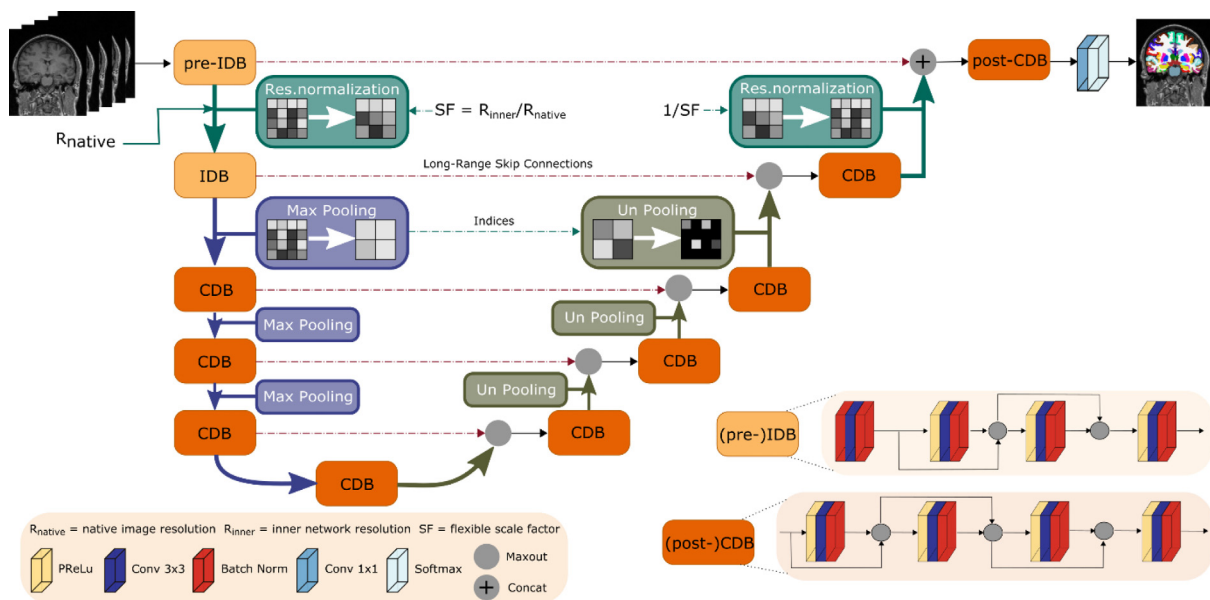


Fig. 3. Voxel size independence in FastSurferVINN. Flexible transitions between resolutions become possible by replacement of (un)pooling with our network-integrated resolution-normalization (green) after the first encoder (pre-IDB) and before the last decoder block (post-CDB). Scale transitions between the other competitive dense blocks (CDB) remain as standard MaxPool and UnPool operations. Each CDB is composed of four sequences of parametric rectified linear unit (PReLU), convolution (Conv) and batch normalization (BN). In the first two encoder blocks ((pre)-IDB), the PReLU is replaced with a BN to normalize the inputs.

ABIDE-II The Autism Brain Imaging Data Exchange II (Di Martino et al., 2017) contains cross-sectional data and focuses on autism spectrum disorders. The dataset contains 1114 subjects from 19 different institutions in total and is accessible online³. The 3D magnetization prepared rapid gradient echo (MPRAGE) sequence, or a vendor specific variant, was used to acquire all data using 3T GE, Philips, and Siemens scanners. Voxel resolutions are not standardized across sites although most scans are acquired at 1.0 mm. In this work we use only the ETH_1 sub-cohort, which provides 0.9 mm HiRes scans acquired on a 3T Philips Achieva with a repetition time (TR) of 3 ms, an echo time (TE) of 3.90 ms, an inversion time (TI) of 1.15 ms, and flip angle of 8°. Since HiRes datasets often exclusively feature Siemens scanners, a subset of 25 HiRes scans from ABIDE-II ETH_1 covering an age range of 20–31 years serves as an independent test set to evaluate generalizability to this unseen Philips scanner and 0.9 mm resolution.

HCP The Human Connectome Project Young Adult (Van Essen et al., 2012) is a cross-sectional study including 3T MRIs from 1200 healthy participants acquired on a customized Siemens scanner (Connectome Skyra). It provides 0.7 mm isotropic de-faced scans of individuals between ages 22 and 35. All scans follow the same MPRAGE protocol with TR 2.4 s, TE 2.14 ms, TI 1 s, and flip angle of 8°. The full dataset is available online.⁴ The Human Connectome Lifespan Pilot Project (Phase 1a)⁵ is an extension of the Young Adult project and contains imaging data from five age groups. Five participants per age group 25–35, 45–55, and 65–75 were scanned on a 3T Siemens Connectome scanner at an isotropic voxel resolution of 0.8 mm following the Young Adult protocol except for a slightly smaller TE of 2.12 ms. In the present study, 30 cases from the Young Adult dataset are used for network training and 20 for validation. A total of 80 cases are used in the final test set. Further, 10 scans from the Lifespan Pilot Project are assembled into a separate test set to assess generalizability to another dataset at 0.8 mm in Section 4.4.

Rhineland Study The Rhineland Study (Breteler et al., 2014) is a large cohort population dataset spanning ages 30 to 90. The 0.8 mm isotropic T1-weighted MRI data is acquired on a 3T Siemens Magnetom Prisma scanner using a multi-echo MPRAGE (ME-MPRAGE) sequence with TR

2560 ms, 10 TEs (1.68 ms, 3.29 ms, 4.90 ms, 6.51 ms, 6 × 5.0 ms), TI 1100 ms, and flip angle 7°. Two separate sets of 30 and 20 subjects are selected for network training and validation, respectively. The final testing set contains another 80 RS subjects.

All datasets were processed using the open source neuroimage analysis suite FreeSurfer⁶ (Fischl (2012); Fischl et al. (2002)). In particular, the FreeSurfer v7.1.1 HiRes stream (Zaretskaya et al., 2018) was used to generate the desired parcellations following the “Desikan–Killiany–Tourville” (DKT) protocol atlas (Desikan et al., 2006; Klein and Tourville, 2012). We follow the same label mapping approach as in FastSurferCNN (Henschel et al., 2020). In short, identical cortical regions on the left and right hemisphere are joined into one class unless they are in close proximity to each other reducing the total number of labels from 95 (DKT without corpus callosum segmentations) to 78 during network training. The affiliation to the left or right hemisphere are restored in the final prediction by estimating the closest WM centroid (left or right hemisphere) to each label cluster. A list of all segmentation labels is provided in the Appendix (see Table 5). In accordance with the FreeSurfer HiRes stream, all MRI brain volumes are conformed to a uniform coordinate orientation at their respective isotropic voxel resolutions and robustly normalized to unsigned characters (0... 255), i.e. the trained network does not depend on skull stripping or bias-field removal.

3.2. Voxel size independent neural network

As mentioned above, inside a standard UNet (Ronneberger et al., 2015), fixed transitions between scales (i.e. resolutions) are performed via down- and up-sampling operations (e.g. pooling and unpooling, see Fig. 3), naturally leading to a reduction in information. Theoretically, any scale-transition operation is replaceable with an alternative sampling strategy as long as it still allows gradients to flow effectively through the network. In FastSurferVINN, we use this concept to enforce voxel size independence by changing the first level transition in the encoder and the last in the decoder to a flexible interpolation step (i.e. network-integrated resolution-normalization, see Fig. 3). Thus, variable

³ http://fcon_1000.projects.nitrc.org/indi/abide/abide_II.html.

⁴ <https://www.humanconnectome.org/study/hcp-young-adult>.

⁵ <https://www.humanconnectome.org/study/hcp-lifespan-pilot>.

⁶ <http://surfer.nmr.mgh.harvard.edu/>.

transitions between resolutions without restrictions to pre-defined fixed voxel sizes become possible, both during training and inference.

3.2.1. Network-integrated resolution-normalization

Similar to spatial transformers (Jaderberg et al., 2015), the interpolation-based scale transition can be divided into two parts: (i) calculation of the sampling coordinates (*grid generator*) and (ii) interpolation operation (*sampler*) to retrieve the spatially transferred output map.

The sampling coordinates are calculated based on the scale factor $SF \in \mathbb{R}^+$ – the quotient of the resolution information of the inner normalized scale $\text{Res}_{\text{inner}}$ (a tuneable hyperparameter set to 1.0 mm throughout our experiments) and the input image $\text{Res}_{\text{native}}$. For optimal interpolation, SF is slightly adjusted to ensure integer feature map dimensions. In the first transition step (Fig. 3, pre-IDB to IDB), the output feature map $V \in \mathbb{R}^{H_{\text{inner}} \times W_{\text{inner}} \times C}$ is produced by sampling the input feature map $U \in \mathbb{R}^{H_{\text{native}} \times W_{\text{native}} \times C}$. In the final transition step (Fig. 3, CDB to post-CDB), this process is reversed effectively by using the inverse scale factor $1/SF$. The interpolation itself is performed by applying a sampling kernel to the input map U to retrieve the value at a particular pixel in the output map V . Different interpolation strategies can be defined based on the selection of the sampling kernel. Theoretically, any kernel with definable (sub-)gradients is applicable. Here, we evaluate the bi-linear, bi-cubic, area, and integer sampling kernels (=NN interpolation). The sampling is identical for each channel, hence, conserving the spatial consistency.

3.2.2. Network architecture modifications

The proposed interpolation strategy can, in general, be included in any CNN equipped with pooling-based scale transitions. Here, due to its success in neuroanatomical segmentation, the principal network design is based on FastSurferCNN (Henschel et al., 2020) – a UNet-type network with a series of four competitive dense blocks (CDB) in the encoder and decoder arm separated by a CDB bottleneck layer. In FastSurferVINN (FastSurferVINN) one additional CDB layer is added (Fig. 3 to each arm, i.e. the pre-IDB and post-CDB).

CDB design In FastSurferCNN (Henschel et al., 2020), a CDB is formed by repetitions of the basic composite function consisting of a 5×5 convolution, followed by batch-normalization (BN) and a probabilistic rectified linear unit (pReLU) activation function. In this work we optimize the architecture design slightly by replacing each 5×5 convolution kernel with two 3×3 kernels (see Fig. 3). This keeps the effective receptive field size within each block identical to FastSurferCNN while reducing parameter load. We implement this change in FastSurferVINN and also, for better comparability, in an updated FastSurferCNN version denoted by: FastSurferCNN*. An ablation study detailing the changes from FastSurferCNN to FastSurferCNN* is included in the appendix (Section A.2). As in FastSurferCNN, feature competition is achieved by using maxout (Goodfellow et al., 2013) instead of concatenations (Jégou et al., 2017) in the local skip connections. In order to guarantee normalized inputs to the maxout activation, the feature map stacking operation is always performed after the BN (see position of maxout in CDB design in Fig. 3).

Pre-IDB The additional encoder block in FastSurferVINN (see Fig. 3, pre-IDB) transfers image intensity information from the native image to the latent space and encodes voxel size-dependent information before the internal interpolation step. In contrast to the described CDB, the raw inputs are normalized by first passing them through a BN-Conv-BN combination before adhering to the original composite function scheme (Conv-BN-pReLU) (see Fig. 3, (pre)-IDB).

Post-CDB Akin to the pre-IDB, an additional CDB block in the decoder is used to merge the non-interpolated feature information returned from the network-integrated resolution-normalization step. Both maps are combined via a concatenation operation and then fed to a standard CDB block (see Fig. 3, (post)-CDB). After the final 1×1 convolution a softmax operation returns the desired class probabilities.

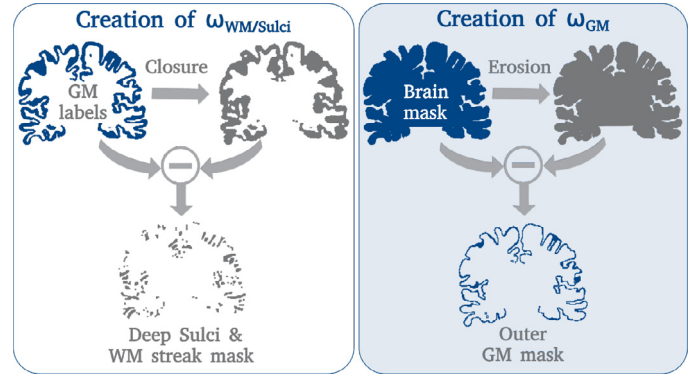


Fig. 4. HiRes weight mask generation. Left: The difference between the gray matter label map (blue) and its closure (gray) produces the deep sulci and WM strand mask. Right: The difference between the original (blue) and eroded (gray) brain mask produces the outer gray matter mask.

3.3. High-res network modifications

To improve segmentation accuracy of detailed structural features, we explore two network modifications, namely a loss-function weighting scheme and an adaptive attention mechanism.

Loss function

The network is trained with a weighted composite loss function of logistic loss and Dice loss (Roy et al., 2017). With $p_{l,i}(x)$ the estimated probability of pixel i to belong to class l and the corresponding ground truth probability y , the loss function can be formulated as

$$\mathcal{L} = - \underbrace{\sum_{l,i} \omega_l y_{l,i} \log p_{l,i}(x)}_{\text{Logistic loss}} - \underbrace{\sum_l \frac{2 \sum_i p_{l,i}(x) y_{l,i}}{\sum_i p_{l,i}(x) + \sum_i y_{l,i}}}_{\text{Soft Dice loss}} \quad (1)$$

with $\omega_l = \omega_{\text{median freq.}} + \omega_{\text{gradient}} + \omega_{\text{GM}} + \omega_{\text{WM/Sulci}}$.

Localized weights Here, $\omega_{\text{median freq.}}$ represents median frequency balancing and ω_{gradient} boundary refinement through a 2D gradient vector (Roy et al., 2017). We now extend ω_l by two weighting terms (ω_{GM} and $\omega_{\text{WM/Sulci}}$) to improve segmentation quality in highly convoluted areas of the cortex that are better represented in submillimeter scans. The WM strand and deep sulci mask ($\omega_{\text{WM/Sulci}}$) emphasizes thin WM strands and narrow sulci, and is defined by the voxels added through a binary closing operation on the gray matter (GM) labels (left side of Fig. 4). The outer gray matter mask (ω_{GM}) accentuates pixels at the boundary of the cortex and is defined by the voxels lost during brain mask erosion (right side of Fig. 4). Overall, ω_{GM} and $\omega_{\text{WM/Sulci}}$ aim to adjust the underlying decision boundary to closely match the target segmentation in PVE-affected locations by assigning higher weights to narrow WM strands, deep sulci, and tissue boundaries with emphasis on the border between cortex and cerebrospinal fluid (CSF).

Adaptive attention module Context-driven learnable attention mechanisms have been used to automatically select optimal scale or filter sizes for specific image regions and can boost segmentation accuracy of differently sized structures within an image (Qin et al., 2018). As a reference, we therefore evaluate the addition of attention in the pre-encoder and post-decoder. The generation of the activation map follows the method proposed in Qin et al. (2018). A detailed description is given in the Appendix (Section A.1). In short, learned activation maps from the attention module are used to dynamically weight each feature response generated by the sequence of convolutions within the CDB. This online weight calculation introduces non-linearities outside the activation function into the CDB.

3.4. View aggregation

In order to account for the inherent 3D geometry of the brain, we adopt the same view aggregation scheme as in Henschel et al. (2020) for

all evaluated models: one F-CNN per anatomical plane is trained and the resulting probability maps are aggregated through a weighted average. Due to missing lateralization in the sagittal view, the number of classes is effectively reduced from 78 to 50 and the weight of the sagittal predictions is reduced by one half compared to the other two views. Inherently, and similar to ensemble learning, the view aggregation combines final soft predictions boosting segmentation accuracy.

3.5. Augmentations

External scaling augmentation for CNNs (exSA) The current state-of-the-art method to introduce resolution invariance into neural networks is extensive scale augmentation (see Fig. 1B). Therefore, we contrast our proposed network-integrated resolution-normalization against this approach. We use random linear transforms with scaling parameters sampled from a uniform distribution of the predefined range 0.8 to 1.15 to augment images during the training phase. Every minibatch hence consist of a potentially scaled MRI (bi-linear interpolation) and a corresponding equally re-scaled label map (NN sampling). To disentangle resolution-independence strategies from micro-architecture changes, we use the same CDB and IDB implementations as for FastSurferVINN described above, with the exception of the resolution-normalization.

Internal scaling augmentation for VINNs (inSA) In order to increase the robustness of the latent space interpolation, we augment the scale factor SF introduced by the VINN with a parameter α , so $SF = \text{Res}_{\text{inner}}/\text{Res}_{\text{native}} + \alpha$ during the network-integrated resolution-normalization. This effectively introduces small resolution variations within the grid sampling procedure. The values for α are sampled from a Gaussian distribution with parameters sigma=0.1 and mean=0. Overall, this modification can be interpreted as an internal scale augmentation randomly resizing the feature maps in the latent space (as opposed to externally augmenting the native images).

3.6. Evaluation metrics

We use the Dice Similarity Coefficient (DSC) and Average Surface Distance (ASD) to compare different network architectures and modifications against each other, and estimate similarity of the predictions to a number of previously unseen scans with respect to FreeSurfer and manual labels as a reference. Both are standard metrics to evaluate segmentation performance.

Here, the DSC [Dice \(1945\)](#); [Sørensen \(1948\)](#) is defined as twice the intersection of ground truth and prediction divided by the sum of their cardinalities and multiplied by 100. A larger DSC represents better overlap between the segmentations with a maximum value of 100 for perfect agreement. The ASD measures the average distance (in mm) between all points $x \in Y, x' \in P$ on the outer surface of the ground truth (Y) and the prediction (P). It is defined as

$$\text{ASD} = \frac{1}{|Y|+|P|} (\sum_{x \in Y} d(x, P) + \sum_{x' \in P} d(x', Y)) \quad (2)$$

with distance $d(x, P) = \min_{x' \in P} \|x - x'\|_2$ representing the minimum of the Euclidean norm. In contrast to the DSC, a smaller ASD indicates better capture of the segmentation boundaries with a value of zero being the minimum (perfect match). Within each section, improvements in segmentation performance are confirmed by statistical testing (Wilcoxon signed-rank test ([Wilcoxon, 1945](#)) after Benjamini-Hochberg correction ([Benjamini and Hochberg, 1995](#)) for multiple testing) referred to as corrected p throughout the paper.

3.7. Training setup

Training dataset Due to the nature of the experiments performed in this paper, the dataset composition varies between individual sections. An overview of the different trainingsets is given in the Appendix ([Table 4](#)). All representative datasets are balanced with regard

to gender and age. Empty slices were filtered from the volumes. All directly compared networks are trained under the same conditions unless stated otherwise. Experimental setups demanding different datasets for training are separated by vertical white lines between bar plots and/or indicated in the figure legend. We additionally discuss this in the corresponding sections. In the first ablative evaluations of FastSurferVINN ([Sections 4.1 and 4.2](#)), 120 representative subjects are selected for training (60 1.0 mm and 60 submillimeter subjects, see [Table 4: Mix](#)), leaving on average 155 single view planes per subject and a total training size of at least 23k images per network. To determine the generalization performance of FastSurferVINN across resolutions, the training set is changed such that the submillimeter scans are of the same resolution and study (see [Table 4: No 0.7 mm and No 0.8 mm](#)). Similarly, the fixed-resolution networks are trained with 60 or 120 0.8 mm scans (see [Table 4: Only 0.8 mm](#)). For the LowRes version 120 1.0 mm scans are used (see [Table 4: Only 1.0 mm](#)). In Big-FastSurferVINN, the 1.0 mm component is extended to 1255 scans while keeping the same number (60) of HiRes scans as in the original training set (see [Table 4: Mix \(Big\)](#)).

Training parameters Independent models for the coronal, axial, and sagittal plane are implemented in PyTorch ([Paszke et al., 2017](#)) and trained for 70 epochs using one NVIDIA V100 GPU with 32 GB RAM. The modified adam optimizer ([Loshchilov and Hutter, 2019](#)) is used with a learning rate set to 0.001. A cosine annealing schedule ([Loshchilov and Hutter, 2017](#)) adapts the learning rate during training where the number of epochs between two warm restarts is initially set to 10 and subsequently increased by a factor of two. The momentum parameter is fixed at 0.95 to compensate for the relatively small mini batch size of 16 images. For maximum fairness, all networks presented within this paper have been trained under equal hardware and hyper-parameter settings.

4. Results

We group the presentation of results into two blocks: 1. ablative architecture improvements to determine the best performing multi-resolution architecture ([Sections 4.1 and 4.2](#)), and 2. performance analysis to comprehensively characterize the advantages of our VINN on a wider variety of datasets, resolutions, scanners, and variations of the training corpus ([Sections 4.3–4.5](#)). Following best practice in data-science, we utilize completely separate datasets during the evaluations: the validation set [Table 4: Validation](#) (for 1.), and various test sets [Table 4: Testing](#) (for 2.). This avoids data-leakage and overfitting, i.e., it ensures that training, architectural design decisions, and final testing cannot influence each other, which could lead to overly optimistic results.

4.1. Scaling augmentation versus network-integrated resolution-normalization in FastSurferVINN

The central contribution of this paper is the design and evaluation of a Voxel-size Independent Neural Network for (sub)millimeter whole brain segmentation. Here, we compare segmentation performance of our FastSurferVINN, which avoids interpolation of label maps, with several approaches that rely on traditional scaling data augmentation. Each subsequent improvement in segmentation performance is confirmed by statistical testing (corrected $p < 0.05$).

Firstly, the original FastSurferCNN CDBs consecutively perform two 5×5 convolution operations followed by a final 1×1 convolution ([Section 3.2.2](#), FastSurferCNN in [Fig. 5](#)). In total, an average DSC of 88.63 and a SD of 0.317 mm is reached for the subcortical structures while the cortical structures average around 87.09 for the DSC and 0.283 mm for the ASD. Optimization of the CDB design (kernel size of 3×3 , [Section 3.2.2](#)) leads to a significant increase in the DSC and reduction of the ASD on both, the subcortical and cortical structures ([Fig. 5](#), FastSurferCNN*). Particularly, the cortical segmentations are improved with an average DSC of 88.01 and a SD of 0.257 mm. Similarly, addition of external scaling augmentation (exSA) to FastSurferCNN* signifi-




Fig. 5. Ablative optimization of FastSurferCNN and comparison to FastSurferVINN. FastSurferCNN (light green) is optimized through a switch to 3×3 kernels (FastSurferCNN*, green). Addition of data augmentation (external scaling augmentation, FastSurferCNN* + exSA, dark green) improves performance further. VINN equipped with internal scaling augmentation (inSA) (FastSurferVINN, orange) outperforms all other models on both subcortical (left) and cortical (right) structures with respect to Dice Similarity Coefficient (DSC, top) and average surface distance (ASD, bottom). Further addition of external scaling augmentation negatively affects performance (VINN + inSA + exSA). Segmentation results with FastSurferVINN are significantly better compared to all other models (corrected $p < 10^{-7}$).

cantly improves segmentation accuracy on the cortical structures (Fig. 5, FastSurferCNN* + exSA). Performance on the subcortical structures is, however, slightly reduced.

Interestingly, VINN, which avoids label map interpolation altogether, shows a positive effect across structures compared to FastSurferCNN* with and without scaling augmentation (Fig. 5, VINN). Specifically, combination of VINN with our internal scaling augmentation (VINN + inSA, referred to as FastSurferVINN) strengthens segmentation performance significantly. The DSC increases to 89.05 for the subcortical and 88.93 for the cortical structures representing the best value across the compared architectures. Similarly, the ASD is significantly reduced with a distance of 0.293 mm for the subcortical and 0.226 mm for the cortical structures (Fig. 5, FastSurferVINN). Finally, addition of external scaling augmentation negatively impacts segmentation performance of VINN (Fig. 5, VINN + inSA + exSA). On average, performance on the subcortical structures is reduced to the level of FastSurferCNN* + exSA, while the cortical structures still represent the second best result for both, DSC and ASD. Overall, FastSurferVINN outperforms all traditional scale augmentation approaches by a significant margin (corrected $p < 10^{-7}$).

As our network-integrated resolution-normalization directly operates on continuous 2D feature maps in the latent space, various sampling kernels can be incorporated. Here, we evaluate the effect of four different interpolation strategies, namely NN, area, bi-cubic, and bi-linear. As visible in Fig. 6, with the exception of NN changing the sampling kernel does not significantly affect segmentation performance. The NN interpolation (Fig. 6, first column) overall reduces performance by 2% for DSC. The change on the ASD is more severe with a decrease in perfor-




Fig. 6. Effect of sampling kernels on network-integrated resolution-normalization. Comparison of nearest-neighbour (NN, purple), area (light violet), bi-cubic (light yellow), and bi-linear (orange) sampling kernels with respect to the Dice Similarity Coefficient (DSC, top) and the average surface distance (ASD, bottom) for subcortical (left) and cortical (right) structures. Segmentation performance with NN is significantly worse than all other interpolation strategies (corrected $p < 10^{-13}$). Area, bi-cubic and bi-linear give equivalent results.

mance by 13% for the subcortical and 7% for the cortical structures. The bi-linear kernel (Fig. 6, last column) shows the best performance overall with a 89.05 DSC and an ASD of 0.293 mm for the subcortical structures and a 88.93 DSC and an ASD of 0.226 mm for the cortical structures. Given that the bi-linear interpolation is computationally favourable to bi-cubic, we henceforth keep it as the sampling kernel of choice.

4.2. HiRes specific adjustments

The higher resolution in submillimeter scans reduces PVEs and offers a potential to optimize segmentation performance (see Fig. 2). Overall, the changes in tissue and border assignment might, however, not be accurately captured during network training as they represent marginal changes in the loss compared to the whole brain volume. In order to focus on the PVE affected regions (specifically small WM strands and deep sulci), we test two different modifications. First, a HiRes loss function (HiRes Loss) up-weights information from the areas in question. Second, attention mechanisms are introduced to enable automatic re-focusing on important information during network training. Here, we equip FastSurferVINN separately with both modifications and evaluate the change in segmentation performance with respect to the DSC and ASD.

Training FastSurferVINN with the new HiRes loss function (see Fig. 7, + HiRes Loss, right bar) significantly improves segmentation performance on the cortical structures (corrected $p < 10^{-5}$), while maintaining high accuracy on the subcortical structures (no significant change consistent with expectations). A final DSC of 89.3 and an ASD of 0.209 mm is achieved for the cortical structures. Further addition of attention, while simultaneously adjusting the number of feature maps to keep the total number of network parameters of FastSurferVINN constant, does not lead to a significant change in segmentation performance (see Fig. 7, + attention, middle bar). Due to the overall significant im-




Fig. 7. Adaptation of the original loss function through addition of HiRes weights focusing on areas strongly effected by PVEs (HiRes Loss, right bar) significantly improves segmentation performance on the cortical structures (compared to FastSurferVINN with original loss, left bar). Addition of attention (middle bar) does not lead to a significant improvement compared to the baseline. Dice Similarity Coefficient (DSC, top) and average surface distance (ASD, bottom) are shown for subcortical (left) and cortical (right) structures. Cortical structures are significantly better segmented with the HiRes Loss (corrected $p < 10^{-5}$). No significant change was detected on the subcortical structures..

provement on DSC and ASD, the HiRes Loss modification is included in all further comparisons.

4.3. Generalizability

After selection of the best architectures for both, VINN and CNN, we now perform a detailed evaluation in a broad selection of test scenarios (see Table 4: Testing). To highlight the advantages of our VINN, we compare FastSurferVINN with FastSurferCNN* + exSA. The latter is improved over the state-of-the-art FastSurferCNN by architectural updates. The difference between FastSurferCNN* + exSA and FastSurferVINN is exclusively the substitution of external scale augmentation (exSA, Section 3.5) with network-integrated resolution-normalization (see Section 3.2.1) and internal scale augmentation (inSA, Section 3.5).

4.3.1. Across datasets

In this section, we evaluate the generalization capabilities of FastSurferVINN in comparison to scale augmentation on the test corpus to establish the performance metrics across multiple datasets.

FastSurferVINN consistently reaches the best DSC and ASD across all nine datasets (Fig. 8, orange bar). All improvements compared to external scale augmentation (dark green bar) are again significant (corrected $p < 0.01$). The best performance is reached for the submillimeter scans from HCP and RS with a subcortical DSC of 88.78 and 90.07 and a cortical DSC of 89.86 and 89.68, respectively. La5c reaches the highest DSC for the 1.0 mm scans (89.51 and 89.07 subcortical and cortical DSC). Similar to the DSC, the ASD is significantly improved for FastSurferVINN compared to the scale augmentation by around 5.2% for the subcortical and 10.5% for the cortical structures. On the cortical

and subcortical structures the best ASD is reached for HCP (0.31 mm, 0.22 mm), RS (0.28 mm, 0.21 mm) and la5c (0.27 mm, 0.22 mm). The biggest improvement with regards to comparability between FastSurferVINN and scale augmentation can be seen for ABIDE-II. Here, the DSC differs by 1.1 and 2.1% for the subcortical and cortical DSC and around 11% for the ASD. This again reflects the better cross-resolution generalization performance of FastSurferVINN already outlined in the previous section. ADNI and OASIS1 also benefit strongly from FastSurferVINN, specifically for the cortical structures. Here, the DSC and ASD are improved by around 1.66% and 12.1%, respectively. Overall, the cortical structures benefit stronger from the internal resolution-normalization in FastSurferVINN.

4.3.2. Unseen resolutions

A core aspect of FastSurferVINN is the implicit compatibility with a variety of resolutions independent of their explicit presence in the training corpus. In order to investigate the generalization capacity of FastSurferVINN in contrast to scale augmentation approaches, we evaluate the inter- and extrapolation capabilities of the trained networks based on the segmentation performance for three resolutions purposefully excluded during network training.

To this end, we specifically drop either (i) all 30 0.8 mm, or (ii) all 30 0.7 mm scans from the training corpus (see Table 4: No 0.7 mm and No 0.8 mm). In order to ensure comparability with respect to the total number of data points and balance between HiRes and LowRes scans, an equal number of subjects from the other respective HiRes datasets are added. In addition, we evaluate performance using the original mixed training set on 25 subjects from ABIDE-II, representing an unseen scanner type (Philips) and resolution (0.9 mm).

As presented in Fig. 9, FastSurferVINN (orange bar) consistently outperforms traditional scale augmentation (dark green bar) across all resolutions (corrected $p < 10^{-4}$). Segmentation performance on the 0.7 mm scans reflects the network's resolution extrapolation capabilities (training corpus consists of 0.8 mm and 1.0 mm scans only). Here, FastSurferVINN reaches a DSC of 86.49 and 87.50 for the subcortical and cortical structures, respectively, representing a significant increase compared to the traditional scale augmentation approach. The improvements on the ASD are even more pronounced with FastSurferVINN reducing the ASD by 4.5% to 0.397 mm on the subcortical and by 19.7% to 0.294 mm on the cortical structures. Comparison of the interpolation capabilities reflected in the 0.8 mm results (training corpus consists of 0.7 mm and 1.0 mm) paint a similar picture. FastSurferVINN reaches the highest DSC and lowest ASD for both subcortical (DSC: 88.75, ASD: 0.316 mm) and cortical structures (DSC of 88.28 and ASD of 0.273 mm). The difference to traditional scale augmentation approaches is again more evident on the cortical structures. Finally, metrics on the 0.9 mm Philips scans are significantly better with FastSurferVINN. A final DSC of 87.85 and 83.68 and an ASD of 0.311 mm and 0.313 mm is reached on the subcortical and cortical structures, respectively.

We face the limitation that no publicly available 3T datasets exists at finer resolutions than 0.7 mm. In order to further explore the extrapolation capabilities of FastSurferVINN in comparison to traditional scale augmentation (FastSurferCNN* + exSA), we therefore evaluate segmentation performance at lower resolutions vastly outside the training range (1.4 mm and 1.6 mm). Note, that FreeSurfer segmentations are not available at lower native resolutions either, as images at resolutions coarser than 1.0 mm are upsampled to a millimeter voxel resolution in an initial conversion step (Fischl, 2012). Therefore, we downsample the high-resolution FreeSurfer segmentations of HiRes test sets (HCP and RS) by a factor of 2 along each axis using majority voting to generate "ground truth". The intensity images used for inference were resampled using cubic interpolation.

In Fig. 10 we illustrate a strong divergence between FastSurferVINN (orange) and FastSurferCNN* + exSA (dark green) as we test with resolutions increasingly outside the training range. With FastSurferVINN, accuracy stays consistently high for both, 1.4 mm and 1.6 mm down-




Fig. 8. Improved generalization performance of FastSurferVINN across nine datasets. FastSurferVINN (orange) outperforms FastSurferCNN* + external scale augmentation (+ exSA, dark green) across subcortical (left) and cortical structures (right) with respect to Dice Similarity Coefficient (DSC, top) and average surface distance (ASD, bottom). Results are consistently better for all datasets (HCP, RS, ABIDE-II, ABIDE-I, ADNI, IXI, LA5C, OASIS1 and OASIS2).




Fig. 9. Improved generalization performance of FastSurferVINN to resolutions not encountered during network training (here training datasets are customized, see Sections 4.3.2 and A.3). FastSurferVINN (orange) outperforms FastSurferCNN* equipped with external scale augmentation (+ exSA, dark green) with respect to Dice Similarity Coefficient (DSC, top) and average surface distance (ASD, bottom) across subcortical (left) and cortical structures (right). Results are significantly better across all resolutions (0.7 mm, 0.8 mm, and 0.9 mm, corrected $p < 10^{-4}$).




Fig. 10. Superior generalization performance of FastSurferVINN to resolutions vastly outside the training domain (1.4 mm, 1.6 mm). FastSurferVINN (orange) outperforms scale-augmentation (FastSurferCNN* + exSA, green) highlighting its extrapolation capabilities. Results are significantly better with respect to Dice Similarity Coefficient (DSC, top) and average surface distance (ASD, bottom) across subcortical (left) and cortical structures (right) (corrected $p < 0.001$ for 1.4 mm and $p < 10^{-13}$ for 1.6 mm).

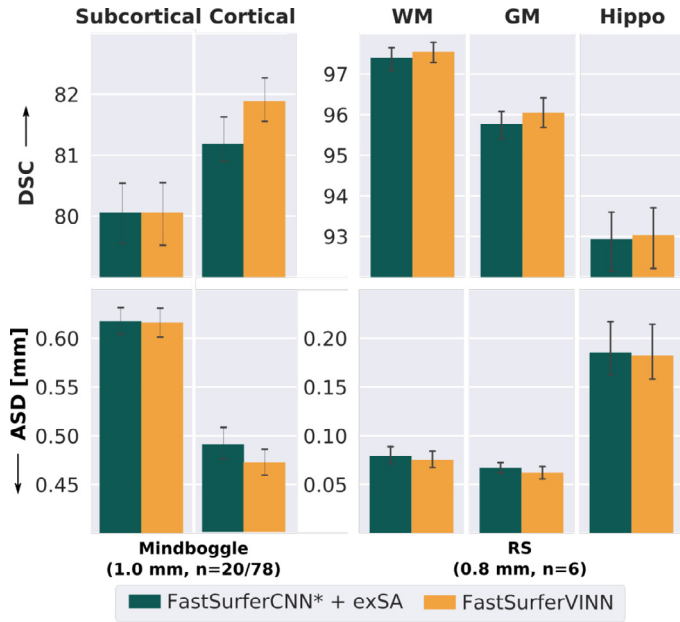


Fig. 11. Performance of FastSurferVINN with respect to manual references. Based on the 1.0 mm scans in Mindboggle101 (left plot) FastSurferVINN (orange) outperforms external scale augmentation (+exSA, dark green) on the cortical structures (right, $N = 78$) with respect to Dice Similarity Coefficient (DSC, top) and average surface distance (ASD, bottom). Results on the subcortical structures (left side, $N = 20$) are equivalent for both approaches. Similarly, segmentation results are better for the 0.8 mm scans of the RS (right plot, $N = 6$) for white matter (WM), gray matter (GM), and hippocampus (Hippo).

sampled images. Here, FastSurferVINN reaches a DSC of above 84.90 and a ASD of 0.473 mm on the subcortical structures (left plot) and a DSC of around 77.70 (77.88 on the 1.4 mm and 77.69 on the 1.6 mm images) and an ASD of 0.480 mm and 0.494 mm on the cortical structures (right plot). The difference to FastSurferCNN* + exSA is again more evident on the cortical structures. Overall, FastSurferVINN outperforms the scaling augmentation by 6.74% on the subcortical and 16.80% on the cortical structures with respect to the DSC when using the downsampled 1.6 mm image as an input. Note, that this comparison is solemnly performed to visualize the performance difference between VINN and augmentation in the absence of high-resolution ground truth beyond the training range. We expect a reduction of performance at low resolutions due to stronger PVE and an 8-fold information reduction.

4.3.3. Comparison to manual reference

Due to the limited availability of manual labels, FreeSurfer segmentations have been used as the reference for comparison so far. In order to account for potential biases we now also evaluate DSC and ASD for the manually edited cortical regions on 78 subjects from the 1.0 mm Mindboggle101 dataset (Klein and Tourville, 2012). The subcortical segmentations available for a subset of 20 subjects within this cohort are used for subcortical evaluations.

On the subcortical structures (see Fig. 11, left part of left plot), FastSurferVINN (orange bar) and traditional scale augmentation (dark green bar) perform equally well with respect to both, DSC (80.26) and ASD (0.616 mm). The difference between the two methods is more pronounced on the cortical structures (Fig. 11, right side of left plot). FastSurferVINN reaches a final DSC of 81.89 and an ASD of 0.471 mm, representing a significant improvement compared to the external scale augmentation approach (corrected $p < 0.005$). These results again indicate that the cortical structures specifically benefit from the internal interpolation approach of FastSurferVINN.

In order to evaluate the benefit of FastSurferVINN for submillimeter scans, we compare segmentation accuracy on three structures (WM, GM and hippocampus) on an in-house set, as no manual HiRes full brain segmentations are publicly available. Note, while we previously reported averages across the 45 cortical and 33 subcortical structures listed in Table 5 for clarity, we now calculate performance measures specifically for these three individual structures. To generate manual annotations, a trained expert corrected FreeSurfer generated segmentations on six cases from the RS. The results are depicted in the left part of Fig. 11. FastSurferVINN again outperforms traditional scale augmentation with a final DSC of 97.54, 96.04, and 93.04 and ASD of 0.075 mm, 0.062 mm, 0.181 mm for WM, GM, and hippocampus, respectively. Consistent with previous results, the GM segmentations show the strongest improvement. Due to the small sample size, no statistical analysis could be performed.

4.4. Fixed-resolution networks versus FastSurferVINN

In order to analyse the advantage of the multi-resolution training approach over a fixed-resolution network, we compare FastSurferVINN with networks trained on (i) 60 0.8 mm subjects, (ii) 120 0.8 mm subjects, and (iii) 120 1.0 mm subjects (see Table 4: Only 0.8 mm and Only 1.0 mm). Inherently, this analysis highlights how the interplay of HiRes information from submillimeter 3T scans and dataset-variations from LowRes MRIs can mutually benefit segmentation performance on the different resolutions. Overall, the analysis of Fig. 12 highlights three important properties.

First, increasing the number of data samples has a strong effect on segmentation performance. The fixed-resolution network trained with 60 0.8 mm scans (left plot; light green bar) reaches an average DSC of 89.28 and ASD of 0.257 mm for RS data (same cohort as in the training set). Doubling the number of cases to 120, significantly improves both measures (corrected $p < 10^{-9}$) with a final DSC of 89.85 and ASD of 0.245 mm (FastSurferVINN, orange bar). Interestingly, FastSurferVINN trained with 60 HiRes (30 0.7 mm and 30 0.8 mm) and 60 LowRes scans (1.0 mm) reaches the same accuracy, as if these scans were actually all HiRes scans from the RS cohort (120 0.8 mm, dark green bar). Benefits from the sample size increase are, therefore, independent of the data resolution. Advantageously, the additional 90 HiRes cases at 0.8 mm may easily be integrated into the FastSurferVINN training dataset, hence, increasing the training corpus with another expected performance gain.

Second, fixed-resolution networks perform well on the cohort they are trained on, but lack generalization capability to other datasets (left plot, HCPL). When comparing networks trained with equally sized datasets (120 HiRes RS fixed-resolution net (dark green bar) and FastSurferVINN (orange bar)) on the 0.8 mm HCPL cohort, FastSurferVINN clearly outperforms the fixed-resolution network with an increase in DSC to 88.99 and a decrease of the ASD to 0.326 mm, representing a significant improvement (corrected $p < 0.005$). The gap in segmentation accuracy between the RS and HCP cohort with respect to DSC is halved for FastSurferVINN compared to the fixed-resolution approach and a similar reduction is visible for the difference in ASD. Third, multi-resolution training benefits 1.0 mm datasets as well. In the right part of Fig. 12, we compare FastSurferVINN (orange bar) and two fixed-resolution 1.0 mm networks – the original FastSurferCNN (CNN, gray bar) as well as the optimized FastSurferCNN* (CNN*, dark green bar) – with DSC (top) and ASD (bottom) metrics. FastSurferCNN is included to highlight the cumulative performance gain of FastSurferVINN (architectural optimization already included in CNN* and voxel-size independence via resolution normalization). The training sets of all three networks contain the exact same subjects, with the only difference being the MRIs resolution: We train the fixed-resolution networks exclusively with 1.0 mm images (native or downsampled from HiRes), while FastSurferVINN uses all images at their native resolution (1.0 mm, 0.8 mm, or 0.7 mm). Note, since the FreeSurfer-based label maps are obtainable




Fig. 12. Flexible- versus fixed-resolution networks. FastSurferVINN (orange) is comparable to, or outperforms, all fixed-resolution networks (green) (left plot: 0.8 mm from RS only, right plots: 1.0 mm) with respect to the Dice Similarity Coefficient (DSC, top) and average surface distance (ASD, bottom). On the submillimeter scans (left plots) generalization to an unseen dataset (HCPL) is significantly improved. Results are consistently better for the 1.0 mm scans (right plot). To highlight cumulative VINN and architectural optimizations, we also compare with the state-of-the-art FastSurferCNN (gray, without optimizations from Sections 4.1 and 4.2, which are already included in CNN*). We retrain this 1 mm fixed-resolution network ensuring equal training datasets and conditions.

at either resolution, we circumvent resampling with NN. Interestingly, the HiRes information from submillimeter scans boosts performance for the 1.0 mm scans. FastSurferVINN reaches a DSC of 87.62 and an ASD of 0.296 mm representing a significant improvement in DSC and ASD compared to the fixed-resolution networks (corrected $p < 0.001$). Compared to the original FastSurferCNN, on average an improvement of 1.46% DSC and 10.06% ASD (0.93 to 2.37% and 8.3 to 12.7%, respectively, across datasets) can be achieved with FastSurferVINN.

Overall, the inherent voxel size independence of FastSurferVINN and resulting multi-resolution training option is highly beneficial to both, submillimeter and 1.0 mm scans.

4.5. Big-FastSurferVINN

Based on the observed performance gain with a larger training corpus, we evaluate whether ASD and DSC can be further improved across resolutions by expanding the training set from 120 to 1315 cases when exclusively adding 1.0 mm scans.

As illustrated in Fig. 13, expansion of the training corpus ($n = 1315$, yellow bar) leads to a noticeable performance gain. Specifically the 1.0 mm dataset benefits from the sample increase with a final DSC of 90.26 and an ASD of 0.231 mm. On average, a 2.95% increase in the DSC and 16.26% decrease in the ASD can be observed for the 1.0 mm scans compared to the smaller training set ($n = 120$, orange bar), representing a significant performance gain (corrected $p < 10^{-20}$). Additionally, ASD and DSC are significantly improved for the 0.7 mm, 0.8 mm and most strongly for the 0.9 mm scans ($p < 10^{-5}$, Wilcoxon signed-rank test). Here, DSC and ASD improve by 3.77% and 24.08%, respectively. Note, that the 0.9 mm dataset is exclusively present in the testing corpus and thus completely new to the networks. A significant resolution bias in the




Fig. 13. Big-FastSurferVINN trained with approximately 20 times more 1.0 mm scans ($n = 1315$, yellow) than the original version ($n = 120$, orange) raises segmentation performance across resolutions. Dice Similarity Coefficient (DSC, top) and average surface distance (ASD, bottom) improve on the submillimeter (0.7 mm–0.9 mm) as well as 1.0 mm scans.

training set, hence, does not decrease, but rather elevate performance across 1.0 mm and submillimeter scans.

5. Discussion

In this paper, we present the first multi-resolution deep learning tool for accurate and efficient (sub)millimeter 3T MRI whole brain segmentation. FastSurferVINN addresses the two main difficulties associated with HiRes MRIs, namely limited availability and diversity specifically for certain groups (e.g. neurodegeneration or scanner types) as well as resolution non-uniformity. Applicability of deep learning approaches is generally restricted to domains where enough training data exists (traditionally 1.0 mm scans). FastSurferVINNs' network-integrated resolution-normalization provides independence to the input voxel grid during both, training and inference, and favourably extends processing spans to resolutions not explicitly included in the training corpus. Further, both, LowRes and HiRes scans, can be included during training and, in turn, benefit from each others favourable properties, i.e. coverage and size versus better representation of detailed structures.

The current state-of-the-art to approximate voxel size independence in networks is data augmentation. Here, transformations with randomly sampled scale parameters are applied to the images (both intensity and label map) in the native space. In FastSurferVINN, we shift this step into the latent space by replacing the first scale transition from a pooling/unpooling operation to interpolation-based down-/upsampling. This network-integrated resolution-normalization interpolates the continuous feature maps on both the encoder and decoder arms and avoids alteration of the underlying ground truth labels entirely.

When optimizing the network architecture, we first show that the best results are achieved with a 3×3 kernel and bilinear interpolation. Also, our HiRes loss which focuses segmentation performance on PVE-affected structures (deep sulci, thin white matter strands and GM-CSF boundary) improves DSC and ASD specifically for the cortical structures. Interestingly, the improvement is consistent across resolutions, indicating that the structural information learned from the HiRes scans is effectively transferred to the 1.0 mm scans. Finally, we evaluate the introduction of attention mechanisms into the architecture. Overall, performance does not improve if the total number of parameters of the network with and without attention is controlled, indicating that network capacity (i.e. number of learnable parameters) rather than adaptive attention is the important factor here.

We demonstrate FastSurferVINN's superior performance compared to state-of-the-art data augmentation for the fast and detailed segmentation of whole brain MRI. Our network-integrated resolution-normalization in combination with internal scale augmentation outperforms traditional scale augmentation in terms of accuracy by a significant margin both with respect to FreeSurfer and a manual standard. Across nine different datasets and four resolutions, including both defaced (HCP) and full-head images, our network achieves the highest DSC compared to FreeSurfer as a reference (88.38 on average), as well as the lowest ASD (0.283 mm on average). In addition, FastSurferVINN achieves the best results on the manually labeled 1.0 mm Mindboggle101 dataset with a DSC of 81.89 and an AVG HD of 0.471 mm on the cortical structures. Correspondingly, WM, GM, and hippocampus are better segmented with FastSurferVINN on six manually corrected scans from the submillimeter RS dataset.

One possible explanation for the consistently improved segmentation metrics with our VINN compared to traditional scale augmentation is the circumvention of label interpolation. For discrete labels, NN or majority voting have to be applied. These kernels are prone to remove important structural details and create jagged segmentation maps (Allebach, 2005; Parker et al., 1983; Schaum, 1993; Thevenaz, 2009). This is underlined by the observation that segmentation performance deteriorates for FastSurferVINN when external scale augmentation is added. Further,

analysis of the interpolation methods during the network-integrated resolution-normalization also highlights the negative impact of the NN interpolation kernels. While all other methods (area, bi-linear, and bi-cubic) performed equally well, performance drops for NN kernels on average by approximately 2% DSC and 10% ASD. Overall, NN seems to have a systematic negative effect on segmentation performance and should, hence, be avoided wherever possible.

In addition, our resolution-normalization interpolates a feature map of vectors instead of a single scalar image slice. As each vector describes a neighborhood in the input image, the available contextual information is extended before the interpolation step and may subsequently support a smoother transition between resolutions. Further, our network-integrated resolution-normalization reduces the range of perceived anatomical size variation in the inner blocks, thereby liberating resources to focus on structural details at a common metric scale (structures have similar sizes after normalization). Finally, a skip-connection transfers the non-interpolated feature maps from the pre-IDB into the post-CDB, potentially improving ambiguous label borders with details available at the native resolution. The differentiation between native voxel-scale and normalized inner scale may specifically be helpful in settings with large inter-subject size similarities (e.g. head size is relatively stable). As such, our network-integrated resolution-normalization approach is not limited to the neuroimaging context, but can be expected to benefit segmentation performance in other domains.

FastSurferVINN further demonstrates excellent generalization performance with respect to inter- and extrapolation, achieving the best results for a variety of unseen submillimeter resolutions excluded from the training corpus. Especially more detailed cortical structures targeted with HiRes acquisition show a significant DSC increase on the 0.7 mm, 0.8 mm and 0.9 mm scans, respectively. Similarly, ASD improved by 19.7%, 10% and 11% highlighting specifically the improved extrapolation capabilities of FastSurferVINN on the left-out 0.7 mm dataset component compared to traditional scale augmentation. Given the lack of a de-facto standard in HiRes datasets, consistent performance towards unseen resolutions is an important property. Specifically, new releases such as the upcoming HCP disease studies⁷ may benefit from a flexible and validated method to process scans with good accuracy at their native submillimeter resolution. Conveniently, FastSurferVINN avoids lossy image downsampling, as well as time intensive re-training, and simplifies re-validation. One future aspect to investigate, is the extent to which this generalization is effective. Unfortunately, with increasing resolutions, evaluation becomes virtually impossible due to (i) unavailability of manual segmentations and (ii) limited support by existing neuroimaging pipelines (i.e. validation is difficult here as well due to point (i)). Our substitute evaluation on the downsampled, lower resolution images underlines that FastSurferVINN robustly generalizes beyond the tested resolution range (see Section 4.3.2, Fig. 10). Specifically, differences to traditional data augmentation are magnified with a more than 16.8% increase in DSC for the 1.6 mm scans. The improved generalizability of FastSurferVINN to under- or unrepresented resolutions especially indicates superior performance in real applications. Furthermore, FastSurferVINN was trained and evaluated on 3T scans only. Adoption to ultra high-field MRI capable of producing resolutions below 0.5 mm is an interesting future avenue, as direct transferability is limited due to the strong signal, contrast, and noise differences (van der Kolk et al., 2013; van der Zwaag et al., 2016).

Overall, FastSurferVINN allows processing of both LowRes and HiRes data without reducing precision on either set. On the contrary, we demonstrated that FastSurferVINN even outperforms a fixed-resolution 1.0 mm network by a significant margin. As the only difference in the training corpus is the increased image resolution of 60 subjects, the gain

⁷ <https://humanconnectome.org/disease-studies>.

in accuracy seems to be predominantly motivated by the structural details provided by the submillimeter scans.

Further, the intensity and demographical variety added by the 1.0 mm scans promoted generalization performance to the HCPL component when FastSurferVINN was compared to a fixed-resolution RS HiRes network. This is of special interest in longitudinal settings, where training on images from only the first time point may introduce a bias towards healthier or younger brains, which can be mitigated by introducing appropriate LowRes out-of-study cases.

Similarly, potential future variations of acquisition parameters within studies may negatively affect performance of fixed-resolution networks and, hence, would require their retraining.

Comparison to fixed-resolution networks revealed another promising application of FastSurferVINN.

As shown in Section 4.4, FastSurferVINN's segmentation performance improves across resolutions even if only additional 1.0 mm scans are included. Surprisingly, performance is dominated by the training corpus size and not significantly influenced by its heterogeneous resolution. On RS, the performance was comparable for FastSurferVINN when training with a mixed resolution corpus (120 total, with only 30 0.8 mm RS cases) as opposed to a fixed-resolution network trained exclusively on 120 0.8 mm RS scans (see Section 4.4). This can be exploited in settings where training cases are scarce at a specific resolution and new acquisitions difficult (e.g. manual segmentations or custom sub-resolution acquisition protocols). FastSurfer-VINN's resolution flexibility is optimally suited to integrate a breadth of already existing resources, reducing the amount of newly acquired data necessary to achieve good performance.

Finally, the results of Big-FastSurferVINN indicate that an imbalance between 1.0 mm and submillimeter data distribution does not introduce a resolution bias.

On the contrary, increasing the 1.0 mm component in Big-FastSurferVINN by a factor 20 improved segmentation performance across all resolutions (0.7 to 1.0 mm). Specifically, the 0.9 mm dataset (ABIDE-II) benefits from the training set extension with an 3.67% increase in DSC and 22.11% decrease in ASD. Unlike HCP and RS, the 0.9 mm ABIDE-II scans were acquired on a Philips scanner, which is present in a larger proportion in the extended training set (30% Philips scanners compared to 1.6%). However, no submillimeter Philips scan was included at any point during training. The improved performance on ABIDE-II, therefore, highlights the potential for FastSurferVINN to actively reduce scanner-biases through inclusion of 1.0 mm MRIs. Exploring an expected alleviation of age- or disease-biases in submillimeter datasets with FastSurferVINN presents an interesting direction for future work.

Overall, we introduce a fast, voxel size independent neural network that scales well to large datasets and enables seamless integration of a variety of resolutions during both, training and inference. Thereby, FastSurferVINN offers the potential to improve generalization performance to future HiRes datasets without retraining, reduce potentially existing dataset biases, and curtail necessary labour and time intensive manual labeling efforts.

FastSurferVINN will be made available as part of the open source FastSurfer (Henschel et al., 2020) package.⁸

Data and code availability statement

All MRI datasets used within this article are publicly available and the weblinks to the open source repositories are cited within the article (3.1 Datasets and Appendix Section A.3. 1.0 mm Datasets) except for the Rhineland Study. The Rhineland Study data is not publicly available because of data protection regulations. 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.

The source code of FastSurferVINN will be made publicly available on Github (<https://github.com/reuter-lab/FastSurfer>) upon acceptance.

Credit authorship contribution statement

Leonie Henschel: Methodology, Software, Validation, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **David Kügler:** Conceptualization, Methodology, Software, Writing – original draft, Writing – review & editing, Supervision. **Martin Reuter:** Conceptualization, Methodology, Resources, Writing – original draft, Writing – review & editing, Supervision, Project administration, Funding acquisition.

Acknowledgments

This work was supported by DZNE institutional funds, by the Federal Ministry of Education and Research of Germany (031L0206, 01GQ1801), and by NIH (R01 LM012719, R01 AG064027, R56 MH121426, and P41 EB030006).

We would like to thank the Rhineland Study group (PI Monique M.B. Breteler) for supporting the data acquisition and management. We also thank Rika Etteldorf for providing corrected high-resolution labels.

Data used in the preparation of this article were obtained in part by the OASIS Cross-Sectional with principal investigators D. Marcus, R. Buckner, J. Csernansky J. Morris; P50 AG05681, P01 AG03991, P01 AG026276, R01 AG021910, P20 MH071616, U24 RR021382, and OASIS: Longitudinal: Principal Investigators: D. Marcus, R. Buckner, J. Csernansky, J. Morris; P50 AG05681, P01 AG03991, P01 AG026276, R01 AG021910, P20 MH071616, U24 RR021382. Further, data used in the preparation of this article were obtained from the MIRIAD database. The MIRIAD investigators did not participate in analysis or writing of this report. The MIRIAD dataset is made available through the support of the UK Alzheimer's Society (Grant RF116). The original data collection was funded through an unrestricted educational grant from GlaxoSmithKline (Grant 6GKC). Data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf. Data collection and sharing for this project was funded by the Alzheimer's Disease Neuroimaging Initiative (ADNI) (National Institutes of Health Grant U01 AG024904) and DOD ADNI (Department of Defense award number W81XWH-12-2-0012). ADNI is funded by the National Institute on Aging, the National Institute of Biomedical Imaging and Bioengineering, and through generous contributions from the following: AbbVie, Alzheimer's Association; Alzheimer's Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; Bristol-Myers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. The Canadian Institutes of Health Research is providing funds to support ADNI clinical sites in Canada. Private sector contributions are facilitated by the Foundation for the National Institutes of Health (www.fnih.org). The grantee organization is the Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Therapeutic Research

⁸ <https://github.com/Deep-MI/FastSurfer>.

Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California. Data were also provided in part by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the [NIH Blueprint for Neuroscience Research](#); and by the McDonnell Center for Systems Neuroscience at Washington University.

Appendix A

A1. Attention

A context-driven learnable attention mechanism following the method proposed in [Qin et al. \(2018\)](#) is evaluated as a reference. This modular mechanism can be inserted into any CDB and offers variable placement across the network architecture. Here, we investigate the addition of attention in the pre-IDB and post-CDB (see [Section 3.2.2](#)) as they are predominantly dealing with feature maps at the native image resolution. In short, a sequence of two convolutions (\tilde{H}_1^a : 3×3 kernel, half the number of filters as the input; \tilde{H}_2^a : 1×1 kernel, one more filter than convolutions in the dense blocks) is followed by a final element-wise softmax function normalizing the activations, denoted as $\lambda_0, \lambda_1, \lambda_2$ and λ_3 , to a sum of 1. The activation map λ_0 is used for weighting the input features representing an effective receptive field of 1. Each voxel is, therefore, assigned a different attention value driven by the image context. In contrast to [Qin et al. \(2018\)](#), the feature responses at different scales (c_1, c_2 and X_a) are generated by a sequence of convolutions within the CDB (H_a^1, H_a^2 and H_a^3). Each convolution increases the receptive field size by $(k_1 - 1 \times k_2 - 1)$ with $(k_1 \times k_2)$ representing the kernel size. The calculated attention maps are then used to weight the local skip connections which combine the outputs of each convolution within these blocks. This is mathematically formulated in [Eqs. \(A.1a\)–\(A.1c\)](#).

$$X_a = \lambda_3 H_a^3(c_2) \quad (\text{A.1a})$$

$$c_2 = \max(\lambda_2 H_a^2(c_1), c_1) \quad (\text{A.1b})$$

$$c_1 = \underbrace{\max(\lambda_1 H_a^1(X_{a-1}), \lambda_0 X_{a-1})}_{\text{Attention CDB}} \quad (\text{A.1c})$$

A2. Ablative optimization of FastSurferCNN

Ablative evaluation of the CNN base architecture optimizations are summarized in [Tables 1](#) and [2](#). Training and validation are based on the multi-resolution datasets listed in [Table 4 \(Mix\)](#). The original FastSurferCNN CDBs consecutively perform two 5×5 convolution operations followed by a final 1×1 convolution ([Section 3.2.2](#), first row in [Tables 1](#) and [2](#)). In total, an average DSC of 88.63 and 87.09 and a ASD of 0.317 mm and 0.283 mm is reached for the subcortical and cortical structures, respectively. Changing the kernel size to 3×3 while keeping the effective receptive field size of 9×9 per CDB constant ([Tables 1](#) and [2](#), $3 \times 3, 64\text{F}$) leads to a significant improvement in DSC ($p < 0.01$,

Table 1
FastSurferCNN optimization. Dice Similarity Coefficient evaluation.

Datasets	DSC	
	Subcortical	Cortical
FastSurferCNN	88.63 (± 1.87)	87.09 (± 2.4)
$+3 \times 3, 64\text{F}$	88.75 (± 1.78)	87.68 (± 2.35)
$+3 \times 3, 71\text{F}$	88.79 (± 1.79)	87.78 (± 2.41)
FastSurferCNN*	88.85 (± 2)	88.01 (± 2.4)

Table 2

FastSurferCNN optimization. Average surface distance evaluation.

Datasets	ASD	
	Subcortical	Cortical
FastSurferCNN	0.317 (± 0.081)	0.283 (± 0.085)
$+3 \times 3, 64\text{F}$	0.312 (± 0.077)	0.267 (± 0.08)
$+3 \times 3, 71\text{F}$	0.311 (± 0.08)	0.264 (± 0.084)
FastSurferCNN*	0.307 (± 0.077)	0.257 (± 0.082)

Wilcoxon-signed rank test). Increasing the number of filters per layer from 64 to 71 preserves the original number of trainable parameters (approximately 1.85×10^6). This change leads to a slight improvement on the subcortical and cortical structures. For comparability to our final VINN, we also add the pre-IDB and post-CDB blocks ([Tables 1](#) and [2](#), FastSurferCNN*). This addition merely assures comparability between the augmentation and interpolation approach. As visible in [Tables 1](#) and [2](#), this change improves segmentation accuracy further with a final DSC of 88.85 and 88.01 and an ASD of 0.307 mm and 0.257 mm. The optimized FastSurferCNN* architecture is used as the baseline for all comparisons (i.e. augmentation and interpolation).

A3. 1.0 mm Datasets

ABIDE I The Autism Brain Imaging Data Exchange I [Di Martino et al. \(2013\)](#) is a cross-sectional study involving 17 international sites and focuses on autism spectrum disorders. It contains data for 1112 individuals between 7 and 64 years of age. Scanner and sequence parameters vary depending on the site and can be accessed on the ABIDE website (https://fcon.1000.projects.nitrc.org/indi/abide/abide_I.html). All MRI data were acquired using 3 Tesla scanners (either Philips, Siemens or GE). 20 cases from the ABIDE-I were used for testing, 68 for training Big-FastSurferVINN.

ADNI The Alzheimer's Disease Neuroimaging Initiative ([Mueller et al., 2005](#)) was launched in 2003 as a public-private partnership, led by principal investigator Michael W. Weiner, MD. The ADNI database has > 2000 participants and is available online at <http://adni.loni.usc.edu>. The primary goal of ADNI has been to test whether serial MRI, positron emission tomography, other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment and early Alzheimer's disease (see www.adni-info.org for up-to-date information). Data were acquired at a resolution of $1.0 \times 1.0 \times 1.2$ mm with 1.5T and 3T-MRIs scanners from the three largest MRI vendors (GE, Philips and Siemens) using an MP-RAGE sequence. Scanner parameters are optimized for the different vendors (see [Jack et al. \(2008\)](#) for details). 15 cases from ADNI where used for training, 8 for validation and 215 for training Big-FastSurferVINN. 40 different cases were used for assessing accuracy and generalizability in the final testset.

IXI This data collection provides 600 MRIs from normal, healthy subjects. The data has been collected at three different sites in London, UK on one GE (1.5T) and two Philips (1.5T and 3T) scanners and is available online (<https://brain-development.org/ixi-dataset/>) under Creative Commons License BY-NC-ND 3.0 (<https://creativecommons.org/licenses/by-sa/3.0/legalcode>). Detailed sequence parameters are available at the given URL. 43 scans from IXI were used for testing and 400 for training Big-FastSurferVINN.

LA5c The cross-sectional UCLA Consortium for Neuropsychiatric Phenomics LA5c Study ([Poldrack et al., 2016](#)) includes 142 individuals diagnosed with a neuropsychiatric or neurodevelopmental disorder (schizophrenia, bipolar disorder, ADHD) and 130 normal controls (ages 21–50). All participants were scanned on a 3T Siemens Trio at a single-center. T1-weighted MP-RAGE images were acquired with field of view of 250, 256x256 matrix, and 176 1.0 mm sagittal partitions.

Table 3

Summary of datasets used for training, validation and testing. Table lists the scanner manufacturer, field strength, (disease) groups, age range and isotropic resolution (Res) used in the paper.

Dataset	Scanner	1.5T/3T	Groups	Age	Res
HCP	Siemens	3T	Normal	22–35	0.7 mm
HCPL	Siemens	3T	Normal	25–75	0.8 mm
RS	Siemens	3T	Normal	30–95	0.8 mm
ABIDE-II ETHZ1	Philips	3T	ASD/Normal	20–31	0.9 mm
ABIDE-I	Philips/GE/Siemens	3T	ASD/Normal	18–64	1.0 mm
ADNI	Philips/GE/Siemens	1.5T/3T	AD/MCI/Normal	55–93	1.0 mm
IXI	Philips/GE	1.5T/3T	Normal	19–87	1.0 mm
LA5C	Siemens	3T	Neuropsych/Normal	21–50	1.0 mm
MBB	Siemens	3T	Normal	20–77	1.0 mm
MIRIAD	GE	1.5T	AD/Normal	55–86	1.0 mm
OASIS1	Siemens	1.5T/3T	Normal	18–90	1.0 mm
OASIS2	Siemens	1.5T/3T	AD/Normal	60–96	1.0 mm

Table 4

Composition of training, validation and testing set used throughout the paper. The asterisk (*) indicates submillimeter datasets are resampled to 1.0 mm.

Usage	Datasets (subjects)	n
Training	Mix	HCP (30), RS (30), ADNI (15), LA5C (16), MIRIAD (7), OASIS1 (14), OASIS2 (8) 120
	No 0.8 mm	HCP (60), ADNI (15), LA5C (16), MIRIAD (7), OASIS1 (14), OASIS2 (8) 120
	No 0.7 mm	RS (60), ADNI (15), LA5C (16), MIRIAD (7), OASIS1 (14), OASIS2 (8) 120
	Only 0.8 mm	RS (small=60, big=120) 60/120
	Only 1.0 mm	HCP* (30), RS* (30), ADNI (15), LA5C (16), MIRIAD (7), OASIS1 (14), OASIS2 (8) 120
	Mix (Big)	HCP (30), RS (30), ABIDE-I (68), ADNI (215), IXI (400), LA5C (203), MBB (195), MIRIAD (30), OASIS1 (79), OASIS2 (65) 1315
	Mix (Big), No 0.8 mm, No 0.7 mm	HCP (20), RS (20), ADNI (8), LA5C (9), MIRIAD (7), OASIS1 (11), OASIS2 (5) 80
	Only 0.8 mm	RS (20) 20
	Only 1.0 mm	HCP* (20), RS* (20), ADNI (8), LA5C (9), MIRIAD (7), OASIS1 (11), OASIS2 (5) 80
	Mix	HCP (80), RS (80), ABIDE-II (25), ABIDE-I (20), ADNI (40), IXI (43), LA5C (15), OASIS1 (30), OASIS2 (17) 350
Validation	No 0.7 mm,	
	No 0.8 mm	HCP (80), RS (80), ABIDE-II (25) 185
	Manual Labels	RS (6), Mindboggle (78) 84
	Only 0.8 mm,	RS (102), HCPL (10), ABIDE-I (20), ADNI (40), IXI (43) 259
	Only 1.0 mm	LA5C (15), OASIS1 (36), OASIS2 (17)
	Mix (Big)	HCP (80), RS (80), ABIDE-II (25), ABIDE-I (20), ADNI (40), IXI (43), LA5C (15), OASIS1 (35), OASIS2 (17) 355

An TI of 1.1 s, TE of 3.5–3.3 ms, TR of 2.53 s and flip angle of 7° was used for all scans. This data was obtained from the OpenfMRI database (<https://openfmri.org/dataset/ds000030/>). Its accession number is ds000030. 16 cases from LA5c were used for training, 9 for validation, 203 for training Big-FastSurferVINN, and 15 for the final testset.

Mindboggle-101 The largest manually corrected set of free, publicly accessible (<https://osf.io/nhtur/>) labeled brain images based on a consistent human cortical labeling protocol (DKTatlas) Klein and Tourville (2012). Mindboggle-101 consists of 101 labeled brain surfaces and volumes derived from T1-weighted brain MRIs of healthy individuals. Except for five subjects (MMRR-3T7T-2, Twins-2, and Afterthought-1), all MRIs are from publicly available collections (i.e. Test-Retest OASIS1 (Marcus et al., 2007), the Multi-Modal Reproducibility Resource (Landman et al., 2011), Nathan Kline Institute Test-Retest and Nathan Kline Institute/Rockland Sample, Human Language Network subjects (Morgan et al., 2009), and Colin Holmes 27 template (Holmes et al., 1998), see (Klein and Tourville, 2012) for details). Manually labeled subcortical segmentations are available for the OASIS1 Test-Retest portion (20 subjects) within Mindboggle-101 (released under Creative Commons License BY-NC-ND 4.0 (<http://creativecommons.org/licenses/by-nc-nd/4.0/>) by Neuromorphometrics, Inc. (<http://Neuromorphometrics.com/>)). All 78 volumes with isotropic voxel sizes are used to evaluate network performance with respect to a manual reference.

MIRIAD The Minimal Interval Resonance Imaging in Alzheimer's Disease (Malone et al., 2013) is a publicly available longitudinal study

with focus on neurodegeneration (see <http://miriad.drc.ion.ucl.ac.uk/>). MRIs of 23 elderly controls and 46 Alzheimer's diseased patients (ages 55+) were acquired at a single center with a 1.5T Signa MRI scanner (GE Medical systems), using an inversion recovery prepared fast spoiled gradient recalled sequence, field of view of 24 cm, 256 × 256 matrix, 124 1.5 mm coronal partitions (voxel size 0.9 × 0.9 × 1.5), TR 15 ms, TE 5.4 ms, flip angle 15°, and TI 650 ms. 7 cases from MIRIAD were used for training and validation, respectively, and 23 for training Big-FastSurferVINN.

Mind-brain-body The MPI Leipzig Mind-Brain-Body cohort (A mind-brain, 2019; Mendes et al., 2019) consists of 321 healthy participants between 20 and 77 years of age. MRIs were acquired on a Siemens Verio 3T with a weighted T1 Magnetization-Prepared 2 Rapid Acquisition Gradient Echoes (MP2RAGE) protocol and sagittal acquisition orientation, one 3D volume with 176 slices, TR of 5000 ms, TE of 2.92 ms, TI1 of 700 ms, TI2 of 2500 ms, flip angle 1 or 4 degrees, flip angle 2 of 5 degrees, echo spacing of 6.9 ms, 1.0 mm isotropic voxel size and field of view of 256 mm. This data was obtained from the OpenfMRI database. Its accession number is ds000221. 195 MRIs were used for training Big-FastSurferVINN.

Oasis-1 Marcus et al.(2007) and Oasis-2 Marcus et al.(2010) The Open Access Series of Imaging Studies 1 and 2, are publicly available (<https://www.oasis-brains.org/>) cross-sectional (Oasis-1) and longitudinal (Oasis-1) studies covering non-demented and demented individuals with very mild to moderate Alzheimer's disease. All subjects were scanned at a single-center using either a 1.5T Vision or a 3T TIM

Table 5
FastSurfer (FastS) internal segmentation IDs and mapping to FreeSurfer (FreeS).

Subcortical structures	FastS	FreeS	Cortical structures	FastS	FreeS
Cerebral white matter (lh)	1	2	caudalanteriorcingulate (lh)	34	1002
Lateral Ventricle (lh)	2	4	caudalmiddlefrontal (lh, rh)	35	1003, 2003
Inferior Lateral Ventricle (lh)	3	5	cuneus (lh)	36	1005
Cerebellar White Matter (lh)	4	7	entorhinal (lh, rh)	37	1006, 2006
Cerebellar Cortex (lh)	5	8	fusiform (lh, rh)	38	1007, 2007
Thalamus (lh)	6	10	inferiorparietal (lh, rh)	39	1008, 2008
Caudate (lh)	7	11	inferiortemporal (lh, rh)	40	1009, 2009
Putamen (lh)	8	12	isthmuscingulate (lh)	41	1010
Pallidum (lh)	9	13	lateraloccipital (lh, rh)	42	1011, 2011
3rd-Ventricle	10	14	lateralorbitofrontal (lh)	43	1012
4th-Ventricle	11	15	lingual (lh)	44	1013
Brain Stem	12	16	medialorbitofrontal (lh)	45	1014
Hippocampus (lh)	13	17	middletemporal (lh, rh)	46	1015, 2015
Amygdala (lh)	14	18	parahippocampal (lh)	47	1016
CSF	15	24	paracentral (lh)	48	1017
Accumbens (lh)	16	26	parsopectacularis (lh, rh)	49	1018, 2018
Ventral DC (lh)	17	28	parsopticalis (lh, rh)	50	1019, 2019
Choroid Plexus (lh)	18	31	parstriangularis (lh, rh)	51	1020, 2020
Cerebral white matter (rh)	19	41	pericalcarine (lh)	52	1021
Lateral Ventricle (rh)	20	43	postcentral (lh)	53	1022
Inferior Lateral Ventricle (rh)	21	44	posteriorcingulate (lh)	54	1023
Cerebellar White Matter (rh)	22	46	precentral (lh)	55	1024
Cerebellar Cortex (rh)	23	47	precuneus (lh)	56	1025
Thalamus (rh)	24	49	rostralanteriorcingulate (lh, rh)	57	1026, 2026
Caudate (rh)	25	50	rostralmiddlefrontal (lh, rh)	58	1027, 2027
Putamen (rh)	26	51	superiorfrontal (lh)	59	1028
Pallidum (rh)	27	52	superiorparietal (lh, rh)	60	1029, 2029
Hippocampus (rh)	28	53	superiortemporal (lh, rh)	61	1030, 2030
Amygdala (rh)	29	54	supramarginal (lh, rh)	62	1031, 2031
Accumbens (rh)	30	58	transversetemporal (lh, rh)	63	1034, 2034
Ventral DC (rh)	31	60	insula (lh, rh)	64	1035, 2035
Choroid Plexus (rh)	32	63	caudalanteriorcingulate (rh)	65	2002
WM-hypointensities	33	77	cuneus (rh)	66	2005
			isthmuscingulate (rh)	67	2010
			lateralorbitofrontal (rh)	68	2012
			lingual (rh)	69	2013
			medialorbitofrontal (rh)	70	2014
			parahippocampal (rh)	71	2016
			paracentral (rh)	72	2017
			pericalcarine (rh)	73	2021
			postcentral (rh)	74	2022
			posteriorcingulate (rh)	75	2023
			precentral (rh)	76	2024
			precuneus (rh)	77	2025
			superiorfrontal (rh)	78	2028

Table 6

Compiled results of Fig. 5 - Ablative optimization of FastSurferCNN and comparison to FastSurferVINN. Microarchitectural changes from FastSurferCNN yield the optimized FastSurferCNN*. Combination of external scale augmentation (+exSA), network-integrated resolution-normalization (VINN) and internal scale augmentation (VINN + inSA = FastSurferVINN) with FastSurferCNN* are shown. Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for the subcortical and cortical structures.

Networks	DSC		ASD	
	Subcortical	Cortical	Subcortical	Cortical
FastSurferCNN	88.63 (\pm 1.87)	87.09 (\pm 2.4)	0.317 (\pm 0.081)	0.283 (\pm 0.085)
FastSurferCNN*	88.85 (\pm 2)	88.01 (\pm 2.4)	0.307 (\pm 0.077)	0.257 (\pm 0.082)
FastSurferCNN* + exSA	88.79 (\pm 1.33)	88.33 (\pm 2.23)	0.311 (\pm 0.066)	0.242 (\pm 0.077)
VINN	88.86 (\pm 1.82)	88.42 (\pm 2.28)	0.308 (\pm 0.081)	0.243 (\pm 0.077)
FastSurferVINN	89.05 (\pm 1.7)	88.93 (\pm 2.01)	0.293 (\pm 0.076)	0.226 (\pm 0.066)
FastSurferVINN + exSA	88.83 (\pm 1.31)	88.6 (\pm 2.13)	0.31 (\pm 0.066)	0.23 (\pm 0.069)

Trio Siemens Scanner in sagittal orientation with a voxel resolution of 1.0x1.0x1.25 mm. For acquisition, a MP-RAGE sequence with a TR of 9.7 ms, TE of 4.0 ms, flip angle of 10° and TI of 20 ms was used. Oasis-1 includes 416 subject between ages 18 to 96, while the longitudinal Oasis-2 focuses on older adults (150 subjects at age 60+). 14 cases from Oasis-1 and 8 from Oasis-2 were used for training, 11 and 5 for valida-

tion, 79 and 65 for training Big-FastSurferVINN and 35 and 17 for final testing.

Participants of the individual studies gave informed consent in accordance with the Institutional Review Board at each of the participating sites. Complete ethic statements are available at the respective study webpages.

Table 7

Compiled results of Fig. 6 - Influence of sampling kernels in the network-integrated resolution-normalization of FastSurferVINN. Comparison of nearest-neighbour (NN), area, bi-cubic and bi-linear sampling kernels. Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for the subcortical and cortical structures.

Networks	DSC		ASD	
	Subcortical	Cortical	Subcortical	Cortical
NN	87.17 (\pm 1.95)	87.24 (\pm 2.3)	0.343 (\pm 0.081)	0.24 (\pm 0.071)
Area	88.99 (\pm 1.72)	88.85 (\pm 2.06)	0.304 (\pm 0.078)	0.228 (\pm 0.068)
Bi-Cubic	89.04 (\pm 1.73)	88.91 (\pm 2.08)	0.304 (\pm 0.079)	0.226 (\pm 0.068)
Bi-Linear	89.05 (\pm 1.7)	88.93 (\pm 2.01)	0.303 (\pm 0.076)	0.226 (\pm 0.066)

Table 8

Compiled results of Fig. 7 - High-Resolution specific adjustments of the original loss function (FastSurferVINN). Addition of attention (+ Attention) or high-resolution weights focusing on areas strongly effected by PVEs (+ HiRes Loss). Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for the subcortical and cortical structures.

Networks	DSC		ASD	
	Subcortical	Cortical	Subcortical	Cortical
FastSurferVINN	89.05 (\pm 1.7)	88.93 (\pm 2.01)	0.303 (\pm 0.076)	0.226 (\pm 0.066)
+ Attention	88.99 (\pm 1.71)	88.79 (\pm 2.03)	0.304 (\pm 0.077)	0.230 (\pm 0.068)
+ HiRes Loss	89.03 (\pm 1.73)	89.30 (\pm 2.07)	0.298 (\pm 0.074)	0.209 (\pm 0.063)

Table 9

Compiled results of Fig. 8: Generalization comparison between optimized FastSurferCNN* + external scale augmentation (exSA) and FastSurferVINN. Dice Similarity Coefficient (DSC) and average surface distance (ASD) for the subcortical and cortical structures across nine different datasets are reported.

		Subcortical		Cortical	
		FastSurferCNN* + exSA	FastSurferVINN	FastSurferCNN* + exSA	FastSurferVINN
DSC	HCP	88.5 (\pm 1.25)	88.78 (\pm 1.2)	89.12 (\pm 1.93)	89.86 (\pm 1.89)
	Rhineland	89.54 (\pm 0.94)	90.07 (\pm 0.88)	88.71 (\pm 2.16)	89.68 (\pm 2.21)
	abide-i	86.55 (\pm 2.37)	86.67 (\pm 2.31)	84.42 (\pm 3.12)	84.97 (\pm 3.13)
	abide-ii	86.9 (\pm 1.2)	87.85 (\pm 1.19)	81.92 (\pm 1.72)	83.68 (\pm 1.55)
	adni	88.65 (\pm 1.86)	88.58 (\pm 2)	87.3 (\pm 2.99)	87.96 (\pm 2.69)
	ixi	86.51 (\pm 3.33)	86.54 (\pm 3.67)	85.94 (\pm 3.01)	86.2 (\pm 2.83)
	la5c	89.49 (\pm 0.51)	89.63 (\pm 0.49)	88.86 (\pm 1.83)	89.46 (\pm 1.49)
	oasis1	88.9 (\pm 0.88)	89.12 (\pm 0.91)	88.33 (\pm 0.88)	89.03 (\pm 0.8)
	oasis2	88.35 (\pm 1.59)	88.49 (\pm 1.65)	87.00 (\pm 2.21)	87.57 (\pm 2.26)
	HCP	0.317 (\pm 0.06)	0.306 (\pm 0.059)	0.245 (\pm 0.126)	0.224 (\pm 0.128)
	Rhineland	0.291 (\pm 0.043)	0.276 (\pm 0.042)	0.241 (\pm 0.13)	0.215 (\pm 0.131)
	abide-i	0.428 (\pm 0.114)	0.421 (\pm 0.108)	0.366 (\pm 0.225)	0.345 (\pm 0.226)
	abide-ii	0.347 (\pm 0.058)	0.311 (\pm 0.054)	0.352 (\pm 0.079)	0.313 (\pm 0.073)
	adni	0.307 (\pm 0.07)	0.31 (\pm 0.075)	0.278 (\pm 0.176)	0.261 (\pm 0.164)
	ixi	0.44 (\pm 0.203)	0.416 (\pm 0.151)	0.296 (\pm 0.107)	0.281 (\pm 0.102)
ASD	la5c	0.279 (\pm 0.038)	0.274 (\pm 0.034)	0.236 (\pm 0.076)	0.217 (\pm 0.062)
	oasis1	0.291 (\pm 0.034)	0.283 (\pm 0.034)	0.233 (\pm 0.031)	0.215 (\pm 0.028)
	oasis2	0.312 (\pm 0.058)	0.309 (\pm 0.065)	0.255 (\pm 0.064)	0.24 (\pm 0.064)

Table 10

Compiled results of Fig. 9 - Generalization of FastSurferCNN* + exSA (external scaling augmentation) and FastSurferVINN to unseen resolutions purposefully excluded from the training set (0.7 mm, 0.8 mm and 0.9 mm). Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for the subcortical and cortical structures. Note, that comparisons between the different data resolutions are not possible due to differing training sets.

Data	Networks	DSC		ASD	
		Subcortical	Cortical	Subcortical	Cortical
0.7 mm	FastSurferCNN* + exSA	85.8 (\pm 1.69)	86.11 (\pm 1.69)	0.416 (\pm 0.088)	0.367 (\pm 0.088)
	FastSurferVINN	86.49 (\pm 1.62)	87.5 (\pm 1.62)	0.397 (\pm 0.084)	0.294 (\pm 0.084)
0.8 mm	FastSurferCNN* + exSA	88.31 (\pm 1.13)	87.31 (\pm 1.13)	0.328 (\pm 0.057)	0.273 (\pm 0.057)
	FastSurferVINN	88.75 (\pm 0.93)	88.28 (\pm 0.93)	0.316 (\pm 0.042)	0.246 (\pm 0.042)
0.9 mm	FastSurferCNN* + exSA	86.9 (\pm 1.2)	81.92 (\pm 1.2)	0.347 (\pm 0.052)	0.352 (\pm 0.052)
	FastSurferVINN	87.85 (\pm 1.19)	83.68 (\pm 1.19)	0.311 (\pm 0.049)	0.313 (\pm 0.049)

Table 11

Compiled results of Fig. 10 - Generalization of FastSurferCNN* + exSA (external scaling augmentation) and FastSurferVINN to resolutions vastly outside the training range (1.4 mm and 1.6 mm). Synthetic ground truth labels and input images have been resampled with majority voting and cubic interpolation from the 0.7 mm and 0.8 mm data to allow comparisons here. Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for the subcortical and cortical structures.

Data	Networks	DSC		ASD	
		Subcortical	Cortical	Subcortical	Cortical
1.4 mm	FastSurferCNN* + exSA	84.54 (\pm 2.12)	76.73 (\pm 2.96)	0.482 (\pm 0.087)	0.597 (\pm 0.198)
	FastSurferVINN	84.94 (\pm 1.65)	77.88 (\pm 2.06)	0.473 (\pm 0.084)	0.480 (\pm 0.136)
	FastSurferCNN* + exSA	80.33 (\pm 6.13)	66.52 (\pm 7.57)	0.758 (\pm 0.329)	1.302 (\pm 0.674)
1.6 mm	FastSurferVINN	85.74 (\pm 1.33)	77.69 (\pm 2.48)	0.473 (\pm 0.062)	0.494 (\pm 0.142)

Table 12

Compiled results of Fig. 11 - Part 1: Performance of FastSurferCNN* + external scale augmentation (exSA) and FastSurferVINN on the 1.0 mm manual reference Mindboggle. Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for the subcortical ($n = 20$) and cortical structures ($n = 78$).

Networks	DSC		ASD	
	Subcortical	Cortical	Subcortical	Cortical
FastSurferCNN* + exSA	80.06 (\pm 1.17)	81.23 (\pm 1.63)	0.617 (\pm 0.03)	0.489 (\pm 0.062)
FastSurferVINN	80.06 (\pm 1.2)	81.89 (\pm 1.67)	0.616 (\pm 0.029)	0.471 (\pm 0.062)

Table 13

Compiled results of Fig. 11 - Part 2: Performance of FastSurferCNN* + external scale augmentation (exSA) and FastSurferVINN on $n = 6$ manually corrected 0.8 mm FreeSurfer labels. Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for the white matter (WM), gray matter (GM) and hippocampus.

Measure		FastSurferCNN* + exSA	FastSurferVINN
DSC	GM	95.77 (\pm 0.46)	96.04 (\pm 0.5)
	WM	97.40 (\pm 0.4)	97.54 (\pm 0.38)
	Hippocampus	92.94 (\pm 1.02)	93.04 (\pm 1.07)
	GM	0.067 (\pm 0.016)	0.062 (\pm 0.018)
	WM	0.079 (\pm 0.025)	0.075 (\pm 0.024)
ASD	Hippocampus	0.185 (\pm 0.078)	0.181 (\pm 0.081)

Table 14

Compiled results of Fig. 12 - Part 1: Comparison of fixed-resolution networks and multi-resolution FastSurferVINN. FastSurferCNN* trained on $n = 60$ or $n = 120$ high-resolution subjects are compared to FastSurferVINN trained on a mixed set (60 high-resolution and 60 1.0 mm scans). Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for the two 0.8 mm datasets (Rhineland study (RS) and Human Connectome Project Lifespan (HCPL)).

Networks	DSC		ASD	
	RS	HCPL	RS	HCPL
FastSurferCNN* ($n = 60$)	89.28 (\pm 0.06)	87.42 (\pm 0.05)	0.391 (\pm 0.05)	0.257 (\pm 0.062)
FastSurferCNN* ($n = 120$)	89.83 (\pm 0.06)	88.14 (\pm 0.05)	0.358 (\pm 0.05)	0.240 (\pm 0.062)
FastSurferVINN ($n = 120$)	89.85 (\pm 0.06)	88.99 (\pm 0.03)	0.326 (\pm 0.033)	0.245 (\pm 0.063)

A4. Tabulated results

In Table 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16. we provide numeric values for DSC and ASD metrics for Figs. 5, 6, 7, 8, 9, 10, 11, 12, 13. A note on reproducibility: Authors wishing to compare their methods to FastSurferVINN and FastSurferCNN are strongly encouraged to download

our code from our repository (github.com/DeepMI/FastSurfer). For reproducibility and fair comparison, we provide a reproducibility guide with training and evaluation instructions for FastSurferVINN and FastSurferCNN on github. Particularly, we recommend re-running evaluations to avoid inconsistencies in non-identical training and evaluation datasets (subjects) as well as reference segmentation.

Table 15

Compiled results of Fig. 12 - Part 2: Comparison of fixed-resolution networks and multi-resolution FastSurferVINN. State-of-the-art FastSurferCNN and optimized FastSurferCNN* trained on 120 1.0 mm scans are compared to FastSurferVINN trained on a mixed set (60 high-resolution and 60 1.0 mm scans). Note, that the scans are equivalent between the training sets and only differ in their resolution. Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for six 1.0 mm datasets.

	Networks	FastSurferCNN ($n = 120$)	FastSurferCNN* ($n = 120$)	FastSurferVINN ($n = 120$)
DSC	ABIDE-I	83.33 (± 4.29)	84.55 (± 3.46)	85.30 (± 2.88)
	ADNI	86.99 (± 2.15)	87.59 (± 2.09)	87.94 (± 1.91)
	IXI	84.46 (± 4.61)	85.22 (± 5.05)	86.14 (± 3.70)
	LA5C	88.3 (± 1.40)	88.84 (± 1.33)	89.25 (± 1.35)
	OASIS1	88 (± 1.38)	88.53 (± 1.37)	88.89 (± 1.39)
	OASIS2	87.36 (± 1.53)	87.81 (± 1.51)	88.18 (± 1.26)
	ABIDE-I	0.441 (± 0.140)	0.416 (± 0.129)	0.402 (± 0.125)
	ADNI	0.315 (± 0.094)	0.294 (± 0.086)	0.281 (± 0.079)
	IXI	0.381 (± 0.128)	0.367 (± 0.148)	0.332 (± 0.106)
	LA5C	0.273 (± 0.045)	0.256 (± 0.044)	0.241 (± 0.043)
ASD	OASIS1	0.278 (± 0.062)	0.264 (± 0.063)	0.255 (± 0.066)
	OASIS2	0.29 (± 0.048)	0.276 (± 0.048)	0.262 (± 0.039)

Table 16

Compiled results of Fig. 13: Influence of training set size. Comparison of FastSurferVINN trained on the original training set ($n = 120$) and a version with approximately 20 times more 1.0 mm scans ($n = 1315$). Dice Similarity Coefficient (DSC) and average surface distance (ASD) are reported for four different resolutions (0.7 mm, 0.8 mm, 0.9 mm and 1.0 mm).

	Datasets	FastSurferVINN + inSA ($n = 120$)	FastSurferVINN + inSA ($n = 1300$)	FastSurferVINN + inSA ($n = 120$)	FastSurferVINN + inSA ($n = 1300$)
DSC	0.7 mm	89.4 (± 1.29)	89.87 (± 1.28)	0.259 (± 0.08)	0.243 (± 0.080)
	0.8 mm	89.84 (± 1.46)	90.57 (± 1.44)	0.241 (± 0.079)	0.218 (± 0.080)
	0.9 mm	85.45 (± 1.13)	88.75 (± 1.39)	0.347 (± 0.050)	0.263 (± 0.052)
	1.0 mm	87.67 (± 2.51)	90.26 (± 1.54)	0.296 (± 0.105)	0.231 (± 0.074)
				ASD	

References

- A mind-brain, 2019. -Body dataset of MRI, EEG, cognition, emotion, and peripheral physiology in young and old adults. *Sci. Data* 6 (1), 180308. doi:10.1038/sdata.2018.308. <http://www.nature.com/articles/sdata2018308>
- Alao, H., Kim, J.-S., Kim, T.S., Lee, K., 2021. Efficient multi-scalable network for single image super resolution. *J. Multimed. Inf. Syst.* 8 (2), 101–110. doi:10.33851/JMIS.2021.8.2.101.
- Allebach, J.P., 2005. 7.1 - image scanning, sampling, and interpolation. In: Bovik, A. (Ed.). *Communications, Networking and Multimedia, Handbook of Image and Video Processing*, second ed.. Academic Press, Burlington doi:10.1016/B978-012119792-6/50115-7. pp. 895–XXVII. <https://www.sciencedirect.com/science/article/pii/B9780121197926501157>
- Bazin, P.-L., Weiss, M., Dinse, J., Schäfer, A., Trampel, R., Turner, R., 2014. A computational framework for ultra-high resolution cortical segmentation at 7 Tesla. *NeuroImage* 93, 201–209. doi:10.1016/j.neuroimage.2013.03.077.
- Benjamini, Y., Hochberg, Y., 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. B* 57 (1), 289–300. doi:10.1111/j.2517-6161.1995.tb02031.x. <https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.2517-6161.1995.tb02031.x>
- Billot, B., Robinson, E., Dalca, A.V., Iglesias, J.E., 2020. Partial volume segmentation of brain MRI scans of any resolution and contrast. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. Springer International Publishing, pp. 177–187. doi:10.1007/978-3-030-59728-3_18.
- 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* 10 (4), P92.
- Chen, H., Dou, Q., Yu, L., Qin, J., Heng, P.-A., 2018. Voxresnet: deep voxelwise residual networks for brain segmentation from 3D MR images. *NeuroImage* 170, 446–455.
- Chen, L.-C., Yang, Y., Wang, J., Xu, W., Yuille, A.L., 2016. Attention to scale: scale-aware semantic image segmentation. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3640–3649. doi:10.1109/CVPR.2016.396. <https://ieeexplore.ieee.org/document/7780765>
- Chen, Z., Zhuang, J., Liang, X., Lin, L., 2019. Blending-target domain adaptation by adversarial meta-adaptation networks. In: *CVPR*, pp. 2248–2257. http://openaccess.thecvf.com/content_CVPR_2019/html/Chen_Blending-Target_Domain_Adaptation_by_Adversarial_Meta-Adaptation_Networks_CVPR_2019_paper.html
- Coupé, P., Mansencal, B., Clément, M., Giraud, R., Denis de Senneville, B., Ta, V.-T., Lepetit, V., Manjon, J.V., 2020. Assemblynet: a large ensemble of CNNs for 3D whole brain MRI segmentation. *NeuroImage* 219, 117026. doi:10.1016/j.neuroimage.2020.117026. <https://www.sciencedirect.com/science/article/pii/S1053811920305127>
- Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., Albert, M.S., Killiany, R.J., 2006. An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *NeuroImage* 31 (3), 968–980.
- Di Martino, A., O'Connor, D., Chen, B., Alaerts, K., Anderson, J.S., Assaf, M., Balsters, J.H., Baxter, L., Beggiato, A., Bernaerts, S., et al., 2017. Enhancing studies of the connectome in autism using the autism brain imaging data exchange II. *Sci. Data* 4, 170010.
- Di Martino, A., Yan, C., Li, Q., Denio, E., Castellanos, F., Alaerts, K., Anderson, J., Assaf, M., Bookheimer, S., Dapretto, M., Deen, B., Delmonte, S., Dinstein, I., Ertl-Wagner, B., Fair, D., Gallagher, L., Kennedy, D., Keown, C., Keyser, C., Lainhart, J., Lord, C., Luna, B., Menon, V., Minshew, N., Monk, C., Mueller, S., Müller, R., Nebel, M., Nigg, J., O'Hearn, K., Pelphrey, K., Peltier, S., Rudie, J., Sunaert, S., Thioux, M., Tyszka, J., Uddin, L., Verhoeven, J., Wenderoth, N., Wiggins, J., Mostofsky, S., Milham, M., 2013. The autism brain imaging data exchange: towards a large-scale evaluation of the intrinsic brain architecture in autism. *Mol. Psychiatry* 19, 659–667.
- Dice, L.R., 1945. Measures of the amount of ecologic association between species. *Ecology* 26 (3), 297–302. doi:10.2307/1932409. <http://doi.wiley.com/10.2307/1932409>
- Dong, C., Loy, C.C., He, K., Tang, X., 2016. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 38 (2), 295–307. doi:10.1109/TPAMI.2015.2439281.
- Fischl, B., 2012. Freesurfer. *NeuroImage* 62 (2), 774–781.
- Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., Van Der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., et al., 2002. Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* 33 (3), 341–355.
- Friston, K., Ashburner, J., Kiebel, S., Nichols, T., Penny, W., 2007. *Statistical Parametric Mapping: The Analysis of Functional Brain Images*. Academic Press.
- Fu, H., Cheng, J., Xu, Y., Wong, D.W.K., Liu, J., Cao, X., 2018. Joint optic disc and cup segmentation based on multi-label deep network and polar transformation. *IEEE Trans. Med. Imaging* 37 (7), 1597–1605. doi:10.1109/TMI.2018.2791488.
- Gaser, C., Dahnke, R., 2016. Cat-a computational anatomy toolbox for the analysis of structural MRI data.
- Gerard, S.E., Herrmann, J., Kaczka, D.W., Musch, G., Fernandez-Bustamante, A., Reinhardt, J.M., 2020. Multi-resolution convolutional neural networks for fully automated segmentation of acutely injured lungs in multiple species. *Med. Image Anal.* 60, 101592. doi:10.1016/j.media.2019.101592. <https://www.sciencedirect.com/science/article/pii/S136184151930132X>
- Gholami, B., Sahu, P., Rudovic, O., Bousmalis, K., Pavlovic, V., 2020. Unsupervised multi-target domain adaptation: an information theoretic approach. *IEEE Trans. Image Process.* 29, 3993–4002. doi:10.1109/TIP.2019.2963389.
- 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., Consortium, W.-M.H., 2013. The minimal preprocessing pipelines for the human connectome project. *NeuroImage* 80, 105–124. doi:10.1016/j.neuroimage.2013.04.127. <https://www.sciencedirect.com/science/article/pii/S1053811913005053>
- Goodfellow, I.J., Warde-Farley, D., Mirza, M., Courville, A., Bengio, Y., 2013. Maxout networks. In: *Proceedings of the 30th International Conference on International Conference on Machine Learning-Volume 28*. JMLR. org, pp. III–1319.

- Gu, F., Burlutskiy, N., Andersson, M., Wilén, L.K., 2018. Multi-resolution networks for semantic segmentation in whole slide images. In: Stoyanov, D., Taylor, Z., Ciompi, F., Xu, Y., Martel, A., Maier-Hein, L., Rajpoot, N., van der Laak, J., Veta, M., McKenna, S., Snead, D., Trucco, E., Garvin, M.K., Chen, X.J., Bogunovic, H. (Eds.), *Computational Pathology and Ophthalmic Medical Image Analysis*. Springer International Publishing, Cham, pp. 11–18.
- Guo, J., Shah, D., Barzilay, R., 2018. Multi-source domain adaptation with mixture of experts. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pp. 4694–4703. doi:10.18653/v1/D18-1498.
- 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. doi:10.1016/j.neuroimage.2020.117012.
- Hoffman, J., Tzeng, E., Park, T., Zhu, J.-Y., Isola, P., Saenko, K., Efros, A., Darrell, T., 2018. CyCADA: cycle-consistent adversarial domain adaptation. In: Dy, J., Krause, A. (Eds.), *Proceedings of the 35th International Conference on Machine Learning*. PMLR, pp. 1989–1998. <https://proceedings.mlr.press/v80/hoffman18a.html>
- Holmes, C.J., Hoge, R., Collins, L., Woods, R., Toga, A.W., Evans, A.C., 1998. Enhancement of MR images using registration for signal averaging. *J. Comput. Assist. Tomogr.* 22, 324–333.
- Huang, L., Zhu, L., Shen, S., Zhang, Q., Zhang, J., 2021. Srnet: scale-aware representation learning network for dense crowd counting. *IEEE Access* 9, 136032–136044. doi:10.1109/ACCESS.2021.3115963.
- Huntenburg, J.M., Steele, C.J., Bazin, P.-L., 2018. Nighres: processing tools for high-resolution neuroimaging. *GigaScience* 7 (7). doi:10.1093/gigascience/giy082.
- Huo, Y., Xu, Z., Xiong, Y., Aboud, K., Parvathaneni, P., Bao, S., Bermudez, C., Resnick, S.M., Cutting, L.E., Landman, B.A., 2019. 3D whole brain segmentation using spatially localized atlas network tiles. *NeuroImage* 194, 105–119.
- Iglesias, J.E., Billot, B., Balbastre, Y., Tabari, A., Conklin, J., Gilberto González, R., Alexander, D.C., Golland, P., Edlow, B.L., Fischl, B., 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. doi:10.1016/j.neuroimage.2021.118206.
- Ito, R., Nakae, K., Hata, J., Okano, H., Ishii, S., 2019. Semi-supervised deep learning of brain tissue segmentation. *Neural Netw.* 116, 25–34.
- Jack, C., A Bernstein, M., C Fox, N., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P., L. Whitwell, J., Ward, C., Dale, A., Felmlee, J., Gunter, J., Hill, D., Kiliany, R., Schuff, N., Fox-Bosetti, S., Lin, C., Studholme, C., Weiner, M., 2008. The alzheimer's disease neuroimaging initiative (ADNI): mri methods. *J. Magn. Reson. Imaging* 27, 685–691.
- Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K., 2015. Spatial transformer networks. In: Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc., <https://proceedings.neurips.cc/paper/2015/file/33ceb07bf44eb3da587e268d663aba1a-Paper.pdf>
- Jégou, S., Drozdal, M., Vazquez, D., Romero, A., Bengio, Y., 2017. The one hundred layers tiramisú: fully convolutional densenets for semantic segmentation. In: *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017 IEEE Conference on. IEEE, pp. 1175–1183.
- Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M., 2012. Fsl. *NeuroImage* 62 (2), 782–790.
- Jin, Y., Wang, X., Long, M., Wang, J., 2020. Minimum class confusion for versatile domain adaptation. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (Eds.), *Computer Vision – ECCV 2020*. Springer International Publishing, Cham, pp. 464–480.
- Kamnitsas, K., Ledig, C., Newcombe, V.F.J., Simpson, J.P., Kane, A.D., Menon, D.K., Rueckert, D., Glocker, B., 2017. Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. *Med. Image Anal.* 36, 61–78. doi:10.1016/j.media.2016.10.004. <https://www.sciencedirect.com/science/article/pii/S1361841516301839>
- Kim, J., Lee, J.K., Lee, K.M., 2016. Accurate image super-resolution using very deep convolutional networks. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1646–1654. doi:10.1109/CVPR.2016.182.
- Kim, J., Lee, J.K., Lee, K.M., 2016. Deeply-recursive convolutional network for image super-resolution. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1637–1645. doi:10.1109/CVPR.2016.181.
- Klein, A., Tourville, J., 2012. 101 labeled brain images and a consistent human cortical labeling protocol. *Front. Neurosci.* 6, 171.
- Landman, B.A., Huang, A.J., Gifford, A., Vikram, D.S., Lim, I.A.L., Farrell, J.A., Bogovic, J.A., Hua, J., Chen, M., Jarso, S., Smith, S.A., Joel, S., Mori, S., Pekar, J.J., Barker, P.B., Prince, J.L., van Zijl, P.C., 2011. Multi-parametric neuroimaging reproducibility: a 3-Tresource study. *NeuroImage* 54 (4), 2854–2866.
- Li, Y., Chen, Y., Wang, N., Zhang, Z.-X., 2019. Scale-aware trident networks for object detection. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 6053–6062. doi:10.1109/ICCV.2019.00615.
- Liu, W., Sun, Y., Ji, Q., 2020. MDan-UNet: multi-scale and dual attention enhanced nested U-Net architecture for segmentation of optical coherence tomography images. *Algorithms* 13 (3). doi:10.3390/a13030060. <https://www.mdpi.com/1999-4893/13/3/60>
- Loshchilov, I., Hutter, F., 2017. SGDR: stochastic gradient descent with warm restarts. In: *International Conference on Learning Representations*. <https://openreview.net/forum?id=Skq89Scxx>
- Loshchilov, I., Hutter, F., 2019. Decoupled weight decay regularization. In: *International Conference on Learning Representations*. <https://openreview.net/forum?id=Bkg6RiQy7>
- Luesebink, F., Wollrab, A., Speck, O., 2013. Cortical thickness determination of the human brain using high resolution 3T and 7T MRI data. *NeuroImage* 70, 122–131. doi:10.1016/j.neuroimage.2012.12.016. <https://www.sciencedirect.com/science/article/pii/S1053811912011937>
- Malone, I.B., Cash, D., Ridgway, G.R., MacManus, D.G., Ourselin, S., Fox, N.C., Schott, J.M., 2013. Miriad-public release of a multiple time point Alzheimer's MR imaging dataset. *NeuroImage* 70, 33–36.
- Marcus, D.S., Fotenos, A.F., Csernansky, J.G., Morris, J.C., Buckner, R.L., 2010. Open access series of imaging studies: longitudinal MRI data in nondemented and demented older adults. *J. Cogn. Neurosci.* 22 (12), 2677–2684.
- Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L., 2007. Open access series of imaging studies (OASIS): cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *J. Cogn. Neurosci.* 19 (9), 1498–1507.
- Markiewicz, C. J., Gorgolewski, K. J., Feingold, F., Blair, R., Halchenko, Y. O., Miller, E., Hardcastle, N., Wexler, J., Esteban, O., Goncalves, M., Jwa, A., Poldrack, R. A., 2021. OpenNeuro: an open resource for sharing of neuroimaging data10.1101/2021.06.28.450168
- McClure, P., Rho, N., Lee, J.A., Kaczmarzyk, J.R., Zheng, C.Y., Ghosh, S.S., Nielson, D.M., Thomas, A.G., Bandettini, P., Pereira, F., 2019. Knowing what you know in brain segmentation using Bayesian deep neural networks. *Front. Neuroinform.* 13, 67. doi:10.3389/fninf.2019.00067. <https://www.frontiersin.org/article/10.3389/fninf.2019.00067>
- Mehta, R., Majumdar, A., Sivaswamy, J., 2017. BrainSegNet: a convolutional neural network architecture for automated segmentation of human brain structures. *J. Med. Imaging* 4 (2), 1–11.
- Mellerio, C., Labeyrie, M.-A., Chassoux, F., Roca, P., Alami, O., Plat, M., Nagara, O., Devaux, B., Meder, J.-F., Oppenheim, C., 2014. 3T MRI improves the detection of transmantle sign in type 2 focal cortical dysplasia. *Epilepsia* 55, 117–122. doi:10.1111/epi.12464. <https://onlinelibrary.wiley.com/doi/full/10.1111/epi.12464>
- Mendes, N., Oligschläger, S., Lauckner, M.E., Golchert, J., Huntenburg, J.M., Falkiewicz, M., Ellamil, M., Krause, S., Baczkowski, B.M., Cozatl, R., Osoianu, A., Kumral, D., Pool, J., Goltz, L., Dreyer, M., Haukeis, P., Jost, R., Kramarenko, Y., Engen, H., Ohrnberger, K., Gorgolewski, K.J., Farrugia, N., Babayan, A., Reiter, A., Schaare, H.L., Reinelt, J., Röbbig, J., Uhlig, M., Erbey, M., Gaebler, M., Smallwood, J., Villringer, A., Margulies, D.S., 2019. A functional connectome phenotyping dataset including cognitive state and personality measures. *Sci. Data* 6 (1), 180307. doi:10.1038/sdata.2018.307. <http://www.nature.com/articles/sdata2018307>
- Morgan, V.L., Mishra, A., Newton, A.T., Gore, J.C., Ding, Z., 2009. Integrating functional and diffusion magnetic resonance imaging for analysis of structure-function relationship in the human language network. *PLoS One* 4 (8), 1–8.
- Mueller, S.G., Weiner, M.W., Thal, L.J., Petersen, R.C., Jack, C.R., Jagust, W., Trojanowski, J.Q., Toga, A.W., Beckett, L., 2005. Ways toward an early diagnosis in Alzheimer's disease: the Alzheimer's disease neuroimaging initiative (ADNI). *Alzheimer's Dementia* 1 (1), 55–66.
- Parker, J.A., Kenyon, R.V., Troxel, D.E., 1983. Comparison of interpolating methods for image resampling. *IEEE Trans. Med. Imaging* 2 (1), 31–39. doi:10.1109/tmi.1983.4307610.
- 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. *NIPS Workshop Autodiff*.
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B., 2019. Moment matching for multi-source domain adaptation. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV). IEEE, Seoul, Korea (South), pp. 1406–1415. doi:10.1109/ICCV.2019.00149. <https://ieeexplore.ieee.org/document/9010750/>
- Poldrack, R., Barch, D., Mitchell, J., Wager, T., Wagner, A., Devlin, J., Cumba, C., Koyejo, O., Milham, M., 2013. Toward open sharing of task-based fMRI data: the openfMRI project. *Front. Neuroinform.* 7, 12. doi:10.3389/fninf.2013.00012. <https://www.frontiersin.org/article/10.3389/fninf.2013.00012>
- Poldrack, R.A., Congdon, E., Triplett, W., Gorgolewski, K., Karlsgodt, K., Mumford, J., Sabb, F., Freimer, N., London, E., Cannon, T., et al., 2016. A phenome-wide examination of neural and cognitive function. *Sci. Data* 3, 160110.
- Poldrack, R.A., Gorgolewski, K.J., 2014. Making big data open: data sharing in neuroimaging. *Nat. Neurosci.* 17 (11), 1510–1517. doi:10.1038/nn.3818.
- Qin, Y., Kamnitsas, K., Ancha, S., Nanavati, J., Cottrell, G.W., Criminisi, A., Nori, A.V., 2018. Autofocus layer for semantic segmentation. In: *Medical Image Computing and Computer Assisted Intervention - MICCAI 2018*, vol. 11072, pp. 603–611. doi:10.1007/978-3-030-00931-1_69. https://link.springer.com/chapter/10.1007/978-3-030-00931-1_69
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer Assisted Intervention - MICCAI 2015*, 9351, pp. 234–241. doi:10.1007/978-3-319-24574-4_28. https://link.springer.com/chapter/10.1007/978-3-319-24574-4_28
- Roy, A.G., Conjeti, S., Navab, N., Wachinger, C., Initiative, A.D.N., et al., 2019. Quicknat: a fully convolutional network for quick and accurate segmentation of neuroanatomy. *NeuroImage* 186, 713–727.
- Roy, A.G., Conjeti, S., Sheet, D., Katouzian, A., Navab, N., Wachinger, C., 2017. Error corrective boosting for learning fully convolutional networks with limited data. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 231–239.
- Roy, S., Krivosheev, E., Zhong, Z., Sebe, N., Ricci, E., 2021. Curriculum graph co-teaching for multi-target domain adaptation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5351–5360.
- Russo, P., Tommasi, T., Caputo, B., 2019. Towards multi-source adaptive semantic segmentation. In: Ricci, E., Rota Bulò, S., Snoek, C., Lanz, O., Messelodi, S., Sebe, N. (Eds.), *Image Analysis and Processing - ICIAP 2019*. Springer International Publishing, Cham, pp. 292–301.

- Schaum, A., 1993. Theory and design of local interpolators. *CVGIP* 55 (6), 464–481. doi:[10.1006/cgip.1993.1035](https://doi.org/10.1006/cgip.1993.1035).
- Shen, J., Wang, Y., Zhang, J., 2021. ASDN: a deep convolutional network for arbitrary scale image super-resolution 26 (1), 13–26. doi:[10.1007/s11036-020-01720-2](https://doi.org/10.1007/s11036-020-01720-2)
- Solano-Castilla, E., Schaefer, A., Reimer, E., Törke, E., Proeger, T., Lohmann, G., Trampel, R., Turner, R., 2011. Parcellation of human amygdala in vivo using ultra high field structural MRI. *NeuroImage* 58, 741–748. doi:[10.1016/j.neuroimage.2011.06.047](https://doi.org/10.1016/j.neuroimage.2011.06.047). <https://www.sciencedirect.com/science/article/pii/S1053811911006926>
- Stankiewicz, J.M., Glanz, B.I., Healy, B.C., Arora, A., Neema, M., Benedict, R.H.B., Guss, Z.D., Tauhid, S., Buckle, G.J., Houtchens, M.K., Khoury, S.J., Weiner, H.L., Guttmann, C.R.G., Bakshi, R., 2011. Brain MRI lesion load at 1.5T and 3T versus clinical status in multiple sclerosis. *J. Neuroimaging* 21, e50–e56. doi:[10.1111/j.1552-6569.2009.00449.x](https://doi.org/10.1111/j.1552-6569.2009.00449.x). <https://onlinelibrary.wiley.com/doi/full/10.1111/j.1552-6569.2009.00449.x>
- Sun, L., Ma, W., Ding, X., Huang, Y., Liang, D., Paisley, J., 2019. A 3D spatially-weighted network for segmentation of brain tissue from MRI. *IEEE Trans. Med. Imaging* 1.
- Sørensen, T.J., 1948. A method of establishing groups of equal amplitude in plant sociology based on similarity of species and its application to analyses of the vegetation on danish commons 5, 1–34.
- Thevenaz, P., 2009. Image interpolation and resampling. In: *Handbook of Medical Image Processing and Analysis*. Elsevier, pp. 465–493. doi:[10.1016/b978-012373904-9.50037-4](https://doi.org/10.1016/b978-012373904-9.50037-4).
- van der Kolk, A.G., Hendrikse, J., Zwanenburg, J.J., Visser, F., Luijten, P.R., 2013. Clinical applications of 7T MRI in the brain. *Eur. J. Radiol.* 82 (5), 708–718. doi:[10.1016/j.ejrad.2011.07.007](https://doi.org/10.1016/j.ejrad.2011.07.007). <https://www.sciencedirect.com/science/article/pii/S0720048X11006450>
- Van Essen, D.C., Ugurbil, K., Auerbach, E., Barch, D., Behrens, T., Bucholz, R., Chang, A., Chen, L., Corbetta, M., Curtiss, S.W., et al., 2012. The human connectome project: a data acquisition perspective. *NeuroImage* 62 (4), 2222–2231.
- van Rijthoven, M., Balkenhol, M., Siliņa, K., van der Laak, J., Ciompi, F., 2021. Hooknet: multi-resolution convolutional neural networks for semantic segmentation in histopathology whole-slide images. *Med. Image Anal.* 68, 101890. doi:[10.1016/j.media.2020.101890](https://doi.org/10.1016/j.media.2020.101890). <https://www.sciencedirect.com/science/article/pii/S1361841520302541>
- Wachinger, C., Reuter, M., Klein, T., 2018. Deepnat: deep convolutional neural network for segmenting neuroanatomy. *NeuroImage* 170, 434–445.
- Wang, H., Yang, W., Lin, Z., Yu, Y., 2019. TMDA: task-specific multi-source domain adaptation via clustering embedded adversarial training. In: 2019 IEEE International Conference on Data Mining (ICDM). IEEE, pp. 1372–1377. doi:[10.1109/ICDM.2019.00176](https://doi.org/10.1109/ICDM.2019.00176).
- Wang, J., Fang, Z., Lang, N., Yuan, H., Su, M.-Y., Baldi, P., 2017. A multi-resolution approach for spinal metastasis detection using deep siamese neural networks. *Comput. Biol. Med.* 84, 137–146. doi:[10.1016/j.combiomed.2017.03.024](https://doi.org/10.1016/j.combiomed.2017.03.024). <https://www.sciencedirect.com/science/article/pii/S0010482517300793>
- Wattjes, M.P., Harzheim, M., Kuhl, C.K., Gieseke, J., Schmidt, S., Klotz, L., Klockgether, T., Schild, H.H., Lutterbey, G.G., 2006. Does high-field MR imaging have an influence on the classification of patients with clinically isolated syndromes according to current diagnostic MR imaging criteria for multiple sclerosis? *Am. J. Neuroradiol.* 27, 1794–1798. <http://www.ajnr.org/content/27/8/1794.long>
- Wilcoxon, F., 1945. Individual comparisons by ranking methods. *Biom. Bull.* 1 (6), 80–83. <http://www.jstor.org/stable/3001968>
- Wolters, A.F., Heijmans, M., Michielse, S., Leentjens, A.F.G., Postma, A.A., Jansen, J.F.A., Ivanov, D., Duits, A.A., Temel, Y., Kuijff, M.L., 2020. The TRACK-PD study: protocol of a longitudinal ultra-high field imaging study in Parkinson's disease. *BMC Neurol.* 20 (1). doi:[10.1186/s12883-020-01874-2](https://doi.org/10.1186/s12883-020-01874-2).
- Xu, R., Chen, Z., Zuo, W., Yan, J., Lin, L., 2018. Deep cocktail network: multi-source unsupervised domain adaptation with category shift. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. IEEE, Salt Lake City, UT, pp. 3964–3973. doi:[10.1109/CVPR.2018.00417](https://doi.org/10.1109/CVPR.2018.00417).
- Xu, Z., Zhang, W., Zhang, T., Li, J., 2021. Hrcnet: high-resolution context extraction network for semantic segmentation of remote sensing images. *Remote Sens.* 13 (1). doi:[10.3390/rs13010071](https://doi.org/10.3390/rs13010071). <https://www.mdpi.com/2072-4292/13/1/71>
- Yang, S., Peng, G., 2018. Attention to refine through multi scales for semantic segmentation. In: *Advances in Multimedia Information Processing – PCM 2018*. Springer International Publishing, pp. 232–241. doi:[10.1007/978-3-030-00767-6_22](https://doi.org/10.1007/978-3-030-00767-6_22).
- Yang, X., Deng, C., Liu, T., Tao, D., 2020. Heterogeneous graph attention network for unsupervised multiple-target domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.* PP.
- Yushkevich, P.A., Pluta, J.B., Wang, H., Xie, L., Ding, S.-L., Gertje, E.C., Mancuso, L., Klöter, D., Das, S.R., Wolk, D.A., 2014. Automated volumetry and regional thickness analysis of hippocampal subfields and medial temporal cortical structures in mild cognitive impairment. *Hum. Brain Mapp.* 36 (1), 258–287. doi:[10.1002/hbm.22627](https://doi.org/10.1002/hbm.22627).
- Zaretskaya, N., Fischl, B., Reuter, M., Renvall, V., Polimeni, J.R., 2018. Advantages of cortical surface reconstruction using submillimeter 7 T MEMPRAGE. *NeuroImage* 165, 11–26. doi:[10.1016/j.neuroimage.2017.09.060](https://doi.org/10.1016/j.neuroimage.2017.09.060).
- Zhang, Y., Brady, M., Smith, S., 2001. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans. Med. Imaging* 20 (1), 45–57. doi:[10.1109/42.906424](https://doi.org/10.1109/42.906424).
- Zhao, S., Li, B., Yue, X., Gu, Y., Xu, P., Hu, R., Chai, H., Keutzer, K., 2019. Multi-source domain adaptation for semantic segmentation. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems*. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2019/file/db9ad56c71619aeed9723314d1456037-Paper.pdf>
- Zheng, Z., Zhang, X., Xiao, P., Li, Z., 2021. Integrating gate and attention modules for high-resolution image semantic segmentation. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* 14, 4530–4546. doi:[10.1109/jstars.2021.3071353](https://doi.org/10.1109/jstars.2021.3071353).
- Zhu, Y., Zhuang, F., Wang, D., 2019. Aligning domain-specific distribution and classifier for cross-domain classification from multiple sources. *Proc. AAAI Conf. Artif. Intell.* 33, 5989–5996. doi:[10.1609/aaai.v33i01.33015989](https://doi.org/10.1609/aaai.v33i01.33015989).
- van der Zwaag, W., Schäfer, A., Marques, J.P., Turner, R., Trampel, R., 2016. Recent applications of UHF-MRI in the study of human brain function and structure: a review. *NMR Biomed.* 29 (9), 1274–1288. doi:[10.1002/nbm.3275](https://doi.org/10.1002/nbm.3275). <https://analyticalsciencejournals.onlinelibrary.wiley.com/doi/pdf/10.1002/nbm.3275>