



Robust machine learning segmentation for large-scale analysis of heterogeneous clinical brain MRI datasets

Benjamin Billot^{a,1} , Colin Magdamo^b, You Cheng^b, Steven E. Arnold^b , Sudeshna Das^b , and Juan Eugenio Iglesias^{a,c,d}

Edited by Jeffrey Ullman, Stanford University (Retired), Stanford, CA; received September 26, 2022; accepted January 24, 2023

Every year, millions of brain MRI scans are acquired in hospitals, which is a figure considerably larger than the size of any research dataset. Therefore, the ability to analyze such scans could transform neuroimaging research. Yet, their potential remains untapped since no automated algorithm is robust enough to cope with the high variability in clinical acquisitions (MR contrasts, resolutions, orientations, artifacts, and subject populations). Here, we present *SynthSeg*⁺, an AI segmentation suite that enables robust analysis of heterogeneous clinical datasets. In addition to whole-brain segmentation, *SynthSeg*⁺ also performs cortical parcellation, intracranial volume estimation, and automated detection of faulty segmentations (mainly caused by scans of very low quality). We demonstrate *SynthSeg*⁺ in seven experiments, including an aging study on 14,000 scans, where it accurately replicates atrophy patterns observed on data of much higher quality. *SynthSeg*⁺ is publicly released as a ready-to-use tool to unlock the potential of quantitative morphometry.

clinical brain MRI | segmentation | deep learning | domain-agnostic

Neuroimaging plays a prominent role in our attempt to understand the human brain, as it enables an array of analyses such as volumetry, morphology, connectivity, physiology, and molecular studies. A prerequisite for almost all these analyses is the contouring of brain structures, a task known as image segmentation. In this context, MRI is the imaging technique of choice, since it enables the acquisition of noninvasive scans in vivo with excellent soft-tissue contrast.

The vast majority of neuroimaging studies rely on prospective datasets of high-quality MRI scans and especially on 1 mm T1-weighted acquisitions. Indeed, these scans present a remarkable white–gray matter contrast and can be easily analyzed with widespread neuroimaging packages, such as SPM (1), FSL (2), or FreeSurfer (3), to derive quantitative morphometric measurements. Meanwhile, brain MRI scans acquired in the clinic (e.g., for diagnostic purposes) present much higher variability in acquisition protocols, and thus cannot be analyzed with conventional neuroimaging softwares. This variability is threefold. First, clinical scans use a wide range of MR sequences and contrasts, which are chosen depending on the tissue properties to highlight. Then, they often present real-life artifacts that are uncommon in research datasets, such as very low signal-to-noise ratio, or incomplete field of view. Finally, instead of using 3D scans at high resolution like in research, physicians usually prefer to acquire a sparse set of 2D images in parallel planes, which are faster to inspect but introduce considerable variability in terms of slice spacing, thickness, and orientation.

The ability to analyze clinical datasets is highly desirable since they represent the overwhelming majority of brain MRI scans. For example, 10 million brain clinical scans were acquired in the US in 2019 alone (4). This figure is orders of magnitude larger than the size of the biggest research studies such as ENIGMA (5) or UK BioBank (6), which comprise tens of thousands of subjects. Hence, analyzing such clinical data would considerably increase the sample size and statistical power of the current neuroimaging studies. Furthermore, it would also enable the analysis of populations that are currently underrepresented in research studies, e.g., UK BioBank and ADNI (7) with 95% white subjects (8, 9), but that are more easily found in clinical datasets. Therefore, there is a clear need for an automated segmentation tool that is robust to MR contrast, resolution, clinical artifacts, and subject populations.

Related Works. The gold standard in brain MRI segmentation is manual delineation. However, this tedious procedure requires costly expertise and is untenable for large-scale clinical applications. Alternatively, one could only consider high-quality scans (i.e., 1 mm T1-weighted scans) that can be analyzed with neuroimaging softwares, but this would

Significance

SynthSeg⁺ is an image segmentation tool for automated analysis of highly heterogeneous brain MRI clinical scans. Our method relies on a new strategy to train deep neural networks, such that it can robustly analyze scans of any contrast and resolution without retraining, which was previously impossible. Moreover, *SynthSeg*⁺ enables scalable quality control of the produced results by automatic detection of faulty segmentations. Our tool is publicly available with FreeSurfer and can be used “out-of-the-box”, which facilitates its use and enhances reproducibility. By unlocking the analysis of heterogeneous clinical data, *SynthSeg*⁺ has the potential to transform neuroimaging studies, given the considerable abundance of clinical scans compared to the size of datasets used in research.

Author contributions: B.B. and J.E.I. designed research; B.B. performed research; B.B. and J.E.I. contributed new reagents/analytic tools; B.B., C.M., Y.C., S.E.A., S.D., and J.E.I. analyzed data; C.M., Y.C., S.E.A., S.D., and J.E.I. edited manuscript; and B.B. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2023 the Author(s). Published by PNAS. This open access article is distributed under [Creative Commons Attribution License 4.0 \(CC BY\)](https://creativecommons.org/licenses/by/4.0/).

¹To whom correspondence may be addressed. Email: benjamin.billot.18@ucl.ac.uk.

This article contains supporting information online at <http://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2216399120/-/DCSupplemental>.

Published February 21, 2023.

drastically decrease effective sample sizes, because such scans are expensive and seldom acquired in the clinic.

Several methods have been proposed for segmentation of MRI scans of variable contrast or resolution. First, contrast-adaptiveness has classically been addressed with Bayesian strategies using unsupervised likelihood model (10). Nevertheless, the accuracy of these methods progressively deteriorates at decreasing resolutions due to partial volume effects, where voxel intensities become less representative of the underlying tissues (11). While such effects can theoretically be modeled within the Bayesian framework (12), the resulting algorithm quickly becomes intractable at decreasing resolutions, thus precluding analysis of clinical scans with large slice thickness.

The modern segmentation literature mostly relies on supervised convolutional neural networks (CNNs) (13, 14), which obtain fast and accurate results on their training domain (i.e., scans with similar contrast and resolution). However, CNNs suffer from the “domain-gap” problem (15), where networks do not generalize well to data with different resolution (16) or MR contrast (17), even within the same modality (e.g., T1-weighted scans acquired with different parameters or hardware) (18). Data augmentation techniques have addressed this problem in intramodality scenarios by applying spatial and intensity transforms to the training data (19). However, the resulting CNNs still need to be retrained for each new MR contrast or resolution, which necessitates costly labeled images. Another approach to bridging the domain gap is domain adaptation, where CNNs are explicitly trained to generalize from a “source” domain with labeled data, to a specific “target” domain, where no labeled examples are available (18, 20). Although these methods alleviate the need for supervision in the target domain, they still

need to be retrained for each new domain, which makes them impractical to apply at scale on highly heterogeneous clinical data.

Very recently, we proposed *SynthSeg* (21), a method that can segment brain scans of any contrast and resolution without retraining. This was achieved by adopting a domain randomization approach (22), where a 3D CNN is trained on synthetic scans of fully randomized contrast and resolution. Consequently, *SynthSeg* learns domain-agnostic representations, which provide it with an outstanding generalization ability compared with previous methods (21). However, *SynthSeg* frequently falters when applied to clinical scans with low signal-to-noise ratio, poor tissue contrast, or acquired at very low resolution—an issue that we address in the present article.

Several strategies have been introduced to improve the robustness of CNNs, most notably hierarchical models. These models divide the final task into easier operations such as progressive refining of segmentations at increasing resolutions (23) or segmenting the same image with increasingly finer labels (24). Although hierarchical models can help improving performance, they may still struggle to produce topologically plausible segmentations in difficult cases, which is a well-known problem for CNNs (25). Recent approaches have sought to solve this problem by modeling high-order topological relations, either by aligning predictions and ground truths in latent space during training (26), by correcting predictions with a registered atlas (27), or with denoising networks (28).

While the aforementioned methods substantially improve robustness, they do not guarantee accurate results in every case. Hence, the ability to identify erroneous predictions is crucial, especially when analyzing clinical data of varying quality.

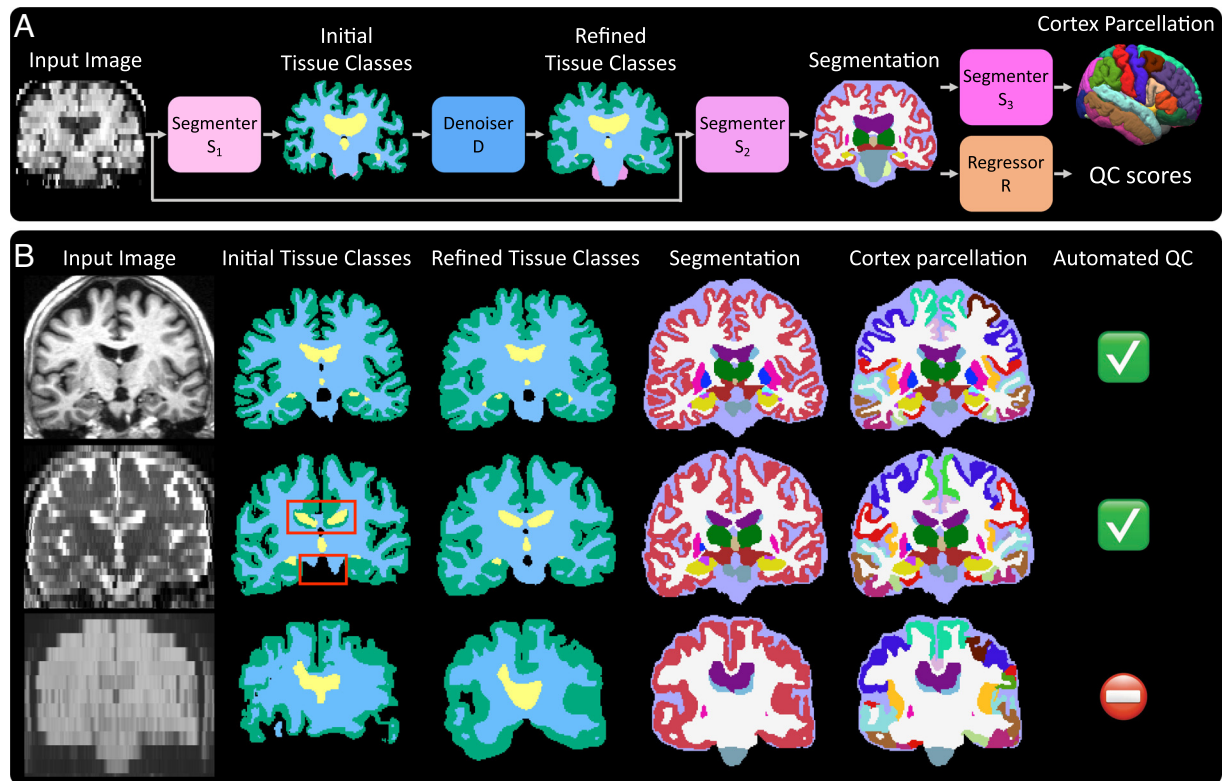


Fig. 1. Overview of *SynthSeg*⁺. (A) Inference pipeline. All modules are implemented as CNNs. (B) Outputs of the intermediate modules for three representative cases. On the first row, all modules obtain accurate results. On the second row, the denoiser corrects mistakes in the initial tissue classes (red boxes), ultimately leading to a good segmentation. Third, the very low tissue contrast of the image leads to a poor segmentation, but the automated QC correctly identifies it as unusable for subsequent analyses.

Traditionally, this has been achieved with visual quality control (QC), but several automated strategies have now been proposed to replace this tedious procedure. A first class of methods seeks to register predictions to a pool of reference segmentations to compute similarity scores (29), but the required registrations remain time-consuming. Therefore, recent techniques employ fast CNNs to model QC as a regression task, where faulty segmentations are rejected by applying a thresholding criterion on the regressed scores (30–32).

Contributions. In this article, we present *SynthSeg*⁺, a clinical brain MRI segmentation suite that is robust to MR contrast, resolution, clinical artifacts, and a wide range of subject populations. Specifically, the proposed method leverages a deep learning architecture composed of hierarchical networks and denoisers. This architecture is trained on synthetic data with the domain randomization approach introduced by *SynthSeg*, and is shown to considerably increase the robustness of the original method to clinical artifacts. Furthermore, *SynthSeg*⁺ includes new modules for cortex parcellation, automated failure detection, and estimation of intracranial volume (ICV, a crucial covariate in volumetry). All aspects of our method are thoroughly evaluated on more than 15,000 highly heterogeneous clinical scans, where *SynthSeg*⁺ is shown to enable automated segmentation and volumetry of large, uncurated clinical datasets. A first version of this work was presented at MICCAI 2022, for whole-brain segmentation only (33). Here, we considerably extend our previous conference paper by adding cortical parcellation, automated QC, and ICV estimation, as well as by evaluating our approach in five new experiments. The proposed method can be run with FreeSurfer using the simple following command:

```
mri_synthseg --i [input] --o [output] --robust
```

Results

A Multitask Segmentation Suite. *SynthSeg*⁺ segments unimodal brain MRI scans by using hierarchical modules designed to efficiently decompose the segmentation task into easier intermediate steps that are less prone to errors (Fig. 1). Specifically, a first network S_1 produces preliminary segmentations of four coarse tissue classes, which are corrected for potential topological

mistakes by a denoiser D (see red boxes in Fig. 1B). Next, predictions of the target labels are obtained with a segmenter S_2 . The obtained cortical region is then further parcellated using a segmenter S_3 . Finally, a regressor R provides us with “QC scores” (describing the quality of the obtained segmentations) for 10 regions, which we use to perform automated QC. We emphasize that all segmentations are given at 1 mm isotropic resolution, regardless of the native resolution of the input.

In the following experiments, we first quantitatively evaluate the accuracy of *SynthSeg*⁺ for whole-brain segmentation and cortex parcellation. Next, we assess the performance of the automated QC module for automatic failure detection. We then present results obtained for ICV estimation, and a proof-of-concept volumetric study for the detection of Alzheimer’s Disease patients. Finally, we demonstrate *SynthSeg*⁺ in a real-life study of aging conducted on more than 14,000 uncurated and heterogeneous clinical scans.

Whole-Brain Segmentation on Clinical Data. In this first experiment, we quantitatively assess the accuracy of *SynthSeg*⁺ for whole-brain segmentation of clinical acquisitions. For this purpose, we use 500 heterogeneous labeled scans, that we take from the picture archiving communication system (PACS) of Massachusetts General Hospital. We compare *SynthSeg*⁺ to *SynthSeg* (21) and three ablations. First, we evaluate an architecture representative of classical cascaded networks ($S_1 + S_2$), by ablating the denoiser D . Then, we test two more variants obtained by appending a denoiser to *SynthSeg* (*SynthSeg*+ D) and the cascaded networks ($S_1 + S_2 + D$), with the aim of evaluating state-of-the-art postprocessing denoisers (28). We measure accuracy with Dice scores, which quantify the overlap between predicted and reference segmentations.

For visualization purposes, we split the results into three classes based on a visual QC performed on the segmentations of *SynthSeg*: “big fails,” “mild fails,” and “passes” (Fig. 2). The results, shown in Fig. 3A, reveal that the hierarchical design of *SynthSeg*⁺ considerably improves robustness (with a mean 76 Dice points) and yields the best scores in all three categories. *SynthSeg*⁺ shows an outstanding improvement of 23.5 Dice points over *SynthSeg* for big fails, and outperforms it by 5.1 and 2.4 points on mild fails and passes, respectively.

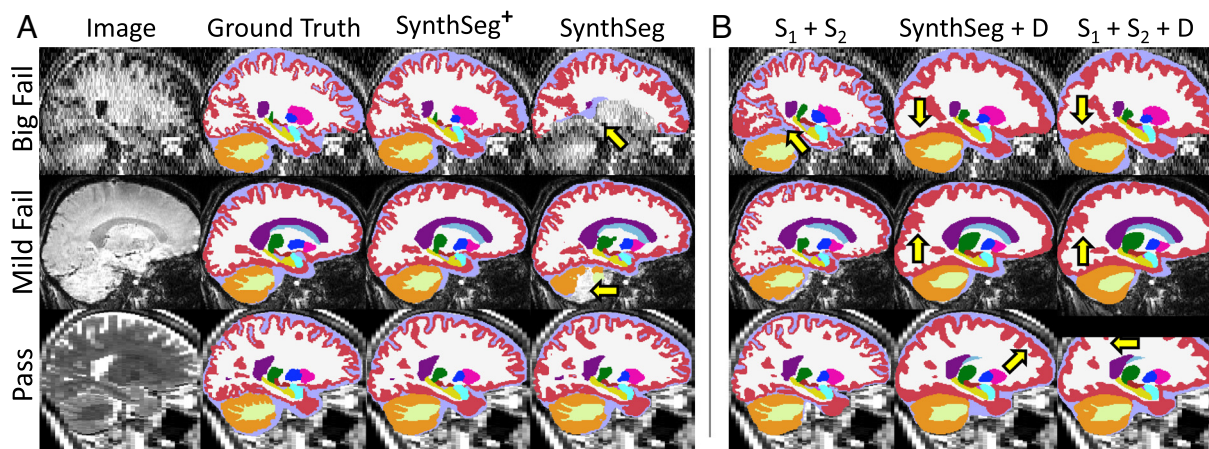


Fig. 2. Segmentations obtained by all tested methods. (A) Comparison between whole-brain segmentations produced by *SynthSeg*⁺ and *SynthSeg*. Here, we show the results obtained for three cases, where *SynthSeg*, respectively, exhibits large (“big fail”), moderate (“mild fail”), and no errors (“pass”). Yellow arrows point at notable mistakes. *SynthSeg*⁺ produces excellent results given the low SNR, poor tissue contrast, or low resolution of the input scans. (B) Segmentations obtained on the same scans by three variants of our method. Note that appending D substantially smooths segmentations.

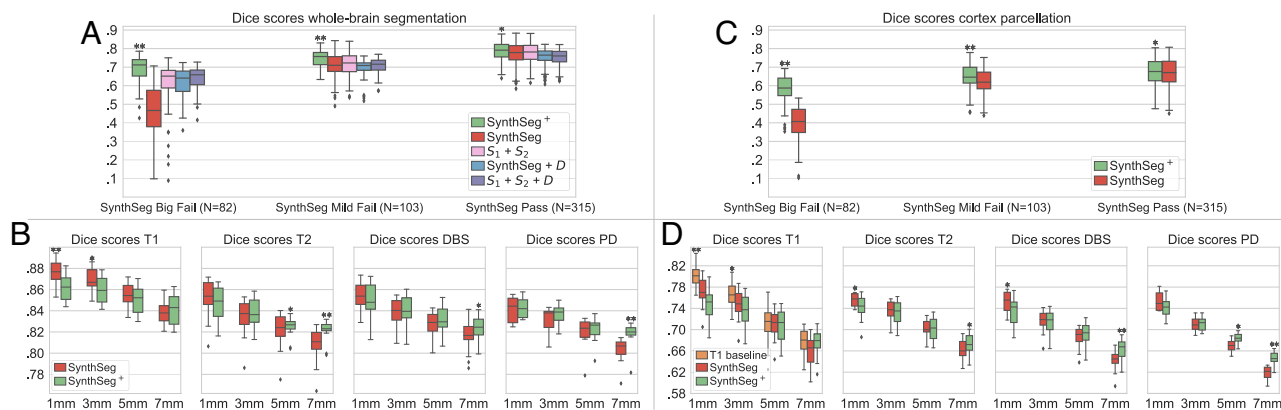


Fig. 3. Dice scores for whole-brain segmentation (A and B) and cortical parcellation (C and D). For (A) and (C), we evaluate the competing methods on 500 heterogeneous clinical scans, presented based on a visual QC of *SynthSeg* segmentations. The results in (B) and (D) are obtained on scans of four MRI modalities at decreasing resolutions. For each dataset, the best method is marked with * or ** if statistically better than the others at a 5% or 1% level (one-sided Bonferroni-corrected Wilcoxon signed-rank test).

In comparison with cascaded networks, employing a denoiser D in *SynthSeg*⁺ to correct the mistakes of S_1 consistently improves the results by 2.3 to 5.1 Dice points. Additionally, using D within our framework (rather than for postprocessing) enables us to substantially outperform the two other variants by 3.1 to 6.2 Dice points. This outcome is explained by the fact that denoisers return very smooth segmentations, especially for the convoluted cortex (Fig. 2). While the denoiser D of our method also exhibits important smoothing effects (Fig. 1), we emphasize that these are successfully recovered by S_2 , which produces segmentations with sharp boundaries.

Robustness Against MR Contrast and Resolution. We now evaluate the performance of *SynthSeg*⁺ as a function of MR contrast and resolution. Here, we use high-resolution research scans (1 mm isotropic resolution) of four MR contrasts, and artificially downsample them at progressively decreasing resolutions. Specifically, we downsample 20 T1-weighted, 18 T2-weighted, 8 proton density (PD), and 18 deep brain stimulation (DBS) scans to nine different resolutions: 3, 5, and 7 mm in axial, coronal, and sagittal orientation.

The results are displayed in Fig. 3B and show that both *SynthSeg* and *SynthSeg*⁺ maintain a high level of accuracy across all tested contrasts and resolutions (above 80 Dice points). The two methods yield very similar scores, but with different trends: While *SynthSeg* produces slightly sharper and more accurate segmentations at 1 mm isotropic resolution (up to 1.4 Dice points better), *SynthSeg*⁺ remains more robust at lower resolutions, where it obtains superior scores for all MRI modalities (maximum gap of 1.6 Dice points).

Cortex Parcellation. Here, we assess the accuracy of *SynthSeg*⁺ for cortex parcellation, and compare it against the results obtained by appending the cortical parcellation module S_3 to *SynthSeg*. Fig. 3C shows that, as in the first experiment, *SynthSeg*⁺ vastly improves the results of *SynthSeg* on heterogeneous clinical data: it yields better scores by 17.1 Dice points on big fails, and is superior by 1.8 points on good cases. Regarding the evolution of performance at decreasing resolutions, Fig. 3D confirms the trend observed for whole-brain segmentation: although *SynthSeg* is more accurate at high resolution (largest gap of 1.6 Dice points for 1 mm T1-weighted scans), it is outperformed by *SynthSeg*⁺ at lower resolutions (2.6 points for 7 mm proton

density scans). We note that the scores for cortex parcellation are below those obtained for whole-brain segmentation, which is not surprising since cortical regions are smaller and more convoluted, and thus harder to segment. Nonetheless, the present results are remarkable since they are similar to the scores obtained by a state-of-the-art supervised CNN trained on T1-weighted scans (19).

Automated Quality Control. We now test the automated QC module of *SynthSeg*⁺ on the 500 clinical scans employed in the previous experiments plus 94 new scans that were almost unusable due to insufficient field of view, wrong organs, critical artifacts, etc. These 594 scans are segmented with *SynthSeg*⁺, and visually classified by an expert neuroanatomist (Y.C.) on a pass/fail basis. This analysis results in a rejection rate of 18.2% for *SynthSeg*⁺, while this rate falls to a remarkable 4.9% when excluding the unusable scans (see examples of fails in *SI Appendix, Appendix 2*). We now seek to automatically replicate the results of this visual QC. Here, *SynthSeg*⁺ is set to reject a segmentation if at least one region obtains an automated QC score below 0.65. We compare *SynthSeg*⁺ against two competitors: a state-of-the-art technique for regression of QC scores (31) and a simpler version of this work, where the segmentation quality is estimated by computing Dice scores between the outputs of D and S_2 (the labels of the latter being converted to the four tissue classes).

Table 1 reports the sensitivity, specificity, accuracy, and area under the ROC curve (AUC) (35) obtained by each method for this binary classification task (ROC curves are given in *SI Appendix*). Despite its relative simplicity, the approach comparing the outputs of D and S_2 already yields a fair accuracy: it correctly classifies 88.4% of the cases, albeit with limited specificity (66.8%). While the two other strategies vastly improve

Table 1. Results of the automated QC analysis on 594 clinical heterogeneous scans for *SynthSeg*⁺ and two competitors

Method	Sensitivity	Specificity	Accuracy	AUC
D vs. S_2 outputs	0.936	0.668	0.884	0.858
Liu et al. (31)	0.982	0.906	0.968	0.989
<i>SynthSeg</i> ⁺ (R)	0.999	0.906	0.981	0.998

The best score for each metric is in bold. No statistical difference is found for the AUC between Liu et al. (31) and our regressor R [$P = 0.081$ when computing a DeLong test (34)].

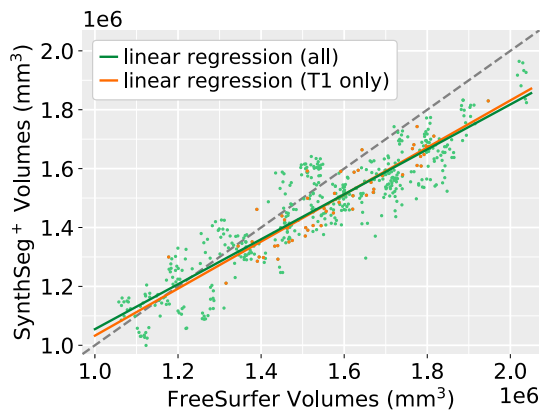


Fig. 4. Scatter plot of ICVs predicted by *SynthSeg*⁺ and FreeSurfer (3) on the 500 clinical scans. Orange points depict 1-mm T1-weighted acquisitions ($N = 62$), while the other scans are in green. The gray dashed line marks where abscissa is equal to ordinate. The Pearson correlation coefficient between the two methods is 0.910 when considering all scans, and 0.906 for T1-weighted scans only ($P < 10^{-9}$ in both cases).

the results (with accuracies above 96%), the simple regressor used in *SynthSeg*⁺ obtains scores very similar to the dedicated state-of-the-art method (31): although no statistical difference can be inferred between AUCs using a DeLong test (34), our approach slightly outperforms Liu et al. for all metrics other than specificity.

Intracranial Volume Estimation. ICV estimation is a crucial task in volumetric studies since it is used to correct head-size effects. Here, we use the clinical dataset of 500 scans to compute correlation coefficients between the ICVs estimated with *SynthSeg*⁺ and FreeSurfer (3). First, we focus on a subset of 62 T1-weighted scans at 1 mm resolution since analyzing these high-quality scans will provide us with a theoretical upper bound for correlation. For these scans, Fig. 4 illustrates that both methods produce strongly correlated ICVs (Pearson's $r = 0.910$). Remarkably, the results hardly change when extending this analysis to all the 500 scans ($r = 0.906$, $p = 0.071$ for a two-sided signed-rank Wilcoxon test between the two ICV distributions), which highlights the robustness of our ICV estimation module against clinical data of varying quality. Finally, a closer inspection of Fig. 4 reveals that *SynthSeg*⁺ tends to predict lower values than FreeSurfer for bigger heads. This observation is consistent with the literature, where similar results have been reported when comparing FreeSurfer to ICVs derived from manual segmentations (36).

Alzheimer's Disease Volumetric Study. We now conduct a proof-of-concept volumetric analysis to study whether *SynthSeg*⁺ can detect subtle hippocampal atrophy linked with Alzheimer's Disease (37). Here, we use a new dataset of 100 subjects, half of whom are diagnosed with Alzheimer's Disease. All subjects are imaged with 1 mm T1-weighted scans, and fluid-attenuated inversion recovery (FLAIR) scans at 5-mm axial resolution, which enables us to study performance across images of varying quality. The predicted hippocampal volumes are linearly corrected for age, gender, and ICV (estimated with *SynthSeg*⁺). Differences between control and diseased populations are then measured by computing effect sizes with Cohen's d (38).

Table 2 reports the results obtained by FreeSurfer, *SynthSeg* and *SynthSeg*⁺. It shows that all methods yield strong effect sizes (between 1.36 and 1.40) on the T1-weighted scans. While segmenting the hippocampus is of modest complexity in 1 mm T1-weighted scans, this task becomes much more challenging on

Table 2. Effect sizes for hippocampal volumes predicted by FreeSurfer, *SynthSeg*, and *SynthSeg*⁺ between 50 controls and 50 Alzheimer's Disease patients

Contrast	Resolution	FreeSurfer	<i>SynthSeg</i>	<i>SynthSeg</i> ⁺
T1	1 mm ³	1.38 (0.895)	1.40 (0.898)	1.36 (0.891)
FLAIR	5 mm axial	-	1.23 (0.876)	1.20 (0.872)

All subjects were imaged with 1-mm T1-weighted scans, as well as 5-mm axial FLAIR scans. AUC scores obtained by every method are shown in parentheses. All approaches produce very similar results, and no statistical difference can be inferred (DeLong tests on AUCs all result in P values above 0.3).

the axial FLAIR scans, where the hippocampus only appears in very few slices (often just 2 to 4). Nonetheless, both *SynthSeg* and *SynthSeg*⁺ maintain a remarkable level of accuracy and are still able to detect strong effect sizes (1.23 and 1.20, respectively). We emphasize that the results obtained are very similar across methods, which is confirmed by the absence of statistical difference when running DeLong tests on corresponding AUCs (all P values are above 0.3).

Aging Study on Over 14,000 Clinical Scans. In this final experiment, we conduct a proof-of-concept clinical volumetric study on 14,752 scans from 1,367 patients with neurology visits at the Massachusetts General Hospital (*Materials and SI Appendix, Appendix 1* for detailed information). Specifically, we verify whether *SynthSeg*⁺ is able to reproduce well-known aging atrophy patterns. We first process all scans with *SynthSeg*⁺ and filter the results with the automated QC module. Instead of rejecting whole segmentations based on a global threshold, we now apply a per-structure rejection criterion, where we only keep the volumes corresponding to regions with a QC score above 0.65. This provides us with between 12,954 (cortex) and 13,357 (white matter) volumes for each region. We then build age-volume trajectories independently for each region with a B-spline model, with linear correction for gender and scan resolution. Remarkably, Fig. 5 shows that the obtained trajectories are highly similar to the results obtained by recent studies on scans of much higher quality (i.e., 1 mm T1-weighted scans) (39–41). For example, *SynthSeg*⁺ accurately replicates the peak in white-matter volume at approximately 30 y of age, the acute increase in ventricular volume for aging subjects, and the early onset of thalamic atrophy compared to the hippocampus and amygdala.

If we now compare it to *SynthSeg*, the proposed approach exhibits similar trajectories but produces far fewer outliers (Fig. 6). This can be seen by the substantially cleaner curves obtained by *SynthSeg*⁺ when including volumes from all available scans, or by the considerably higher number of segmentations that passed the automated QC. Finally, as opposed to *SynthSeg*, *SynthSeg*⁺ produces almost identical results when we evaluate it exclusively on scans of very low resolution (i.e., with slice spacing above 6.5 mm), which further highlights its robustness against acquisitions of widely varying quality.

Discussion

Here, we present *SynthSeg*⁺, a segmentation suite for large-scale analysis of highly heterogeneous clinical brain MRI scans. The proposed method leverages a deep learning architecture, where state-of-the-art domain-agnostic CNNs perform the target segmentation task in a hierarchical fashion. As a result, *SynthSeg*⁺ exhibits an unparalleled robustness to clinical artifacts and can

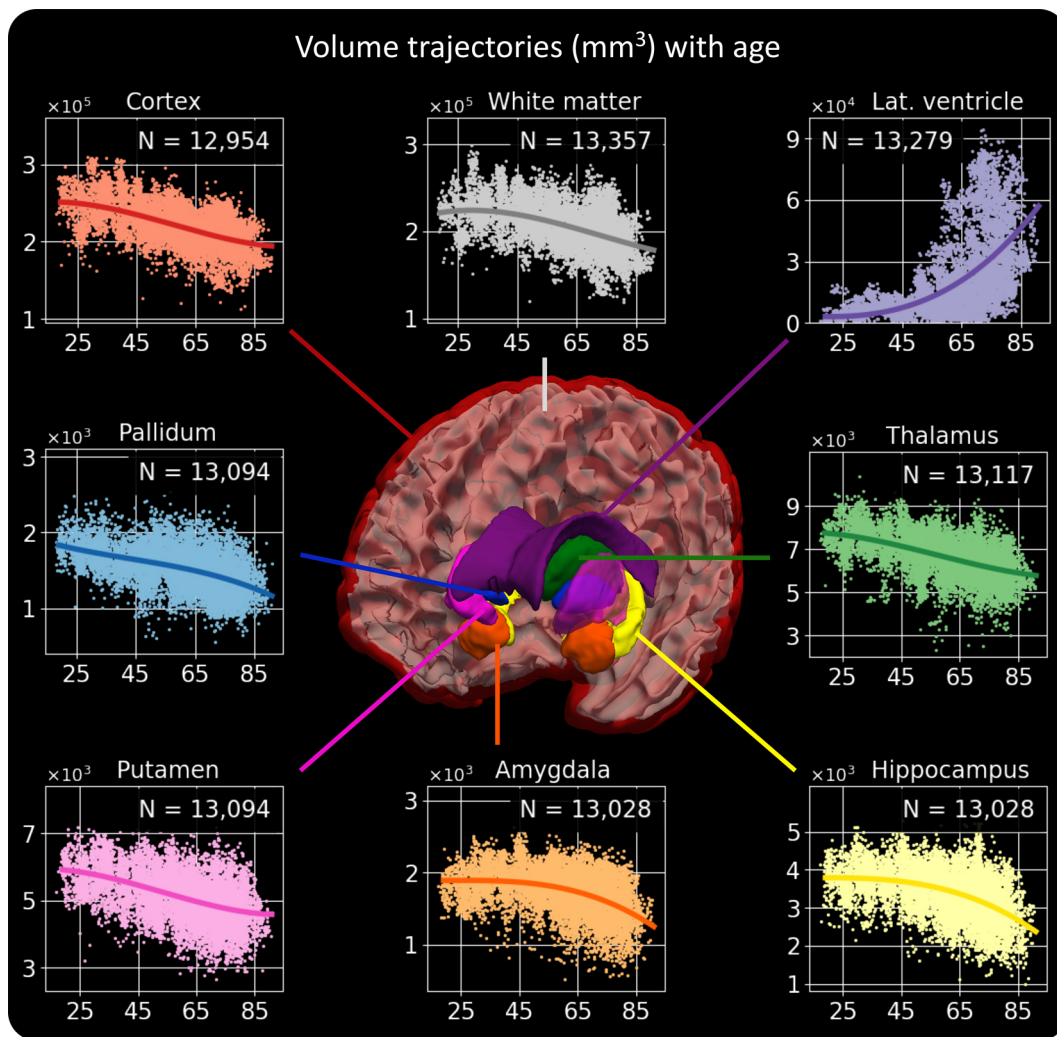


Fig. 5. Volume trajectories obtained by processing 14,752 heterogeneous clinical scans with *SynthSeg*⁺. For each brain region, the value of *N* indicates the number of volumes considered to build the plot, i.e., the number of segmentations that passed the automated QC for this structure. We emphasize that the obtained results are remarkably similar to recent studies, which exclusively employed scans of much higher quality (39–41).

accurately segment clinical data of any MR contrast and resolution, including scans with poor tissue contrast and low signal-to-noise ratio. Moreover, the two proof-of-concept volumetric studies and the quantitative aging experiment (*SI Appendix, Appendix 4*) have demonstrated the robustness of the proposed method against a wide range of subject populations. Finally, our highly precise QC enables us to automatically discard the few erroneous segmentations from subsequent analyses.

The performance of our hierarchical modules is demonstrated throughout all experiments, where it considerably improves the robustness of *SynthSeg* to clinical artifacts, both quantitatively (higher Dice scores in the first experiment), and qualitatively (far fewer outliers in the aging study). Further ablation studies show that using a denoiser *D* to correct the mistakes of *S*₁ leads to a consistent improvement over cascaded networks. Moreover, *SynthSeg*⁺ substantially outperforms state-of-the-art denoising networks for postprocessing (28), which is explained by two reasons. First, these methods do not have access to the input scans and may thus produce segmentations that deviate from the original anatomy. In contrast, we predict final segmentations by exploiting both the input test scans and prior information given by *D*. Second, denoiser predictions are often excessively smooth. Here, we mitigate this issue by using the outputs of *D* as robust priors for *S*₂, which effectively learns to refine

the smooth boundaries given by the denoiser. However, Fig. 2 illustrates that the predictions of *SynthSeg*⁺ may still exhibit slight smoothing effects. While the resolution studies show that this residual smoothness leads to a marginally lower accuracy than *SynthSeg* for scans at high resolution (relatively uncommon in clinical settings), we consider this a minor limitation compared to the considerable gain in robustness of the proposed approach.

SynthSeg⁺ is also able to perform volumetric cortex parcellation of clinical scans in the wild: only FreeSurfer (3) and FastSurfer (42) could tackle this problem automatically, but they only apply to 1 mm T1-weighted scans. The high accuracy of our parcellation module is demonstrated by evaluating it on high-quality scans, for which it yields competitive scores with state-of-the-art supervised CNNs, either tested on the same scans, like the T1-baseline (19) or on different datasets (42). Remarkably, *SynthSeg*⁺ maintains this high level of performance across all tested domains, including the highly heterogeneous clinical acquisitions. We emphasize that this outstanding generalization ability can difficultly be matched by supervised CNNs, which would need to be retrained on every single domain using costly training labeled data.

The results have also shown that the proposed regression-based QC strategy accurately detects the few cases where *SynthSeg*⁺ fails, which are mainly due to acquisitions of poor quality

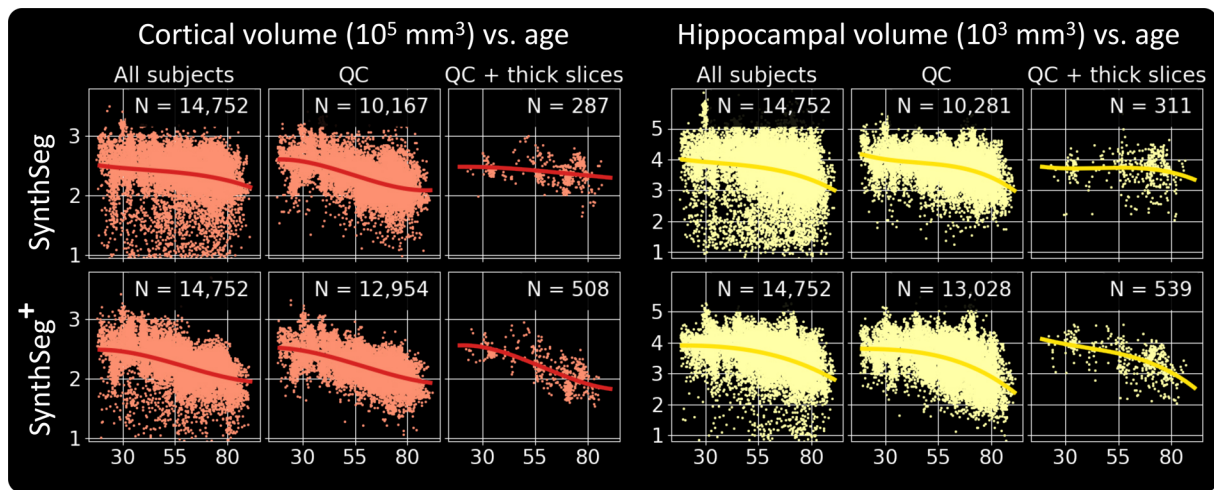


Fig. 6. Cortical and hippocampal volume trajectories obtained by *SynthSeg* and *SynthSeg*⁺ in three different scenarios: using all available scans, keeping only those which passed the automated QC and simulating the case where we only have access to scans acquired at low resolution (here with a slice thickness of more than 6.5 mm). We highlight that *SynthSeg*⁺ is much more robust than *SynthSeg* since it outputs far fewer outliers, and accurately detects atrophy patterns even for scans at very low resolutions.

(e.g., very low SNR, or insufficient coverage of the brain). Interestingly, our method produces slightly better results than the tested state-of-the-art approach (31), despite the fact that the latter was shown to outperform strategies based on direct regression in a simpler problem (3D segmentation with only one label). This outcome can be explained by the fact that variational autoencoders (which are at the core of ref. 31) have been shown in our experiments to produce excessive smoothness when dealing with whole-brain segmentations. We also emphasize that other methods for automated QC were not evaluated in this work since they rely on lengthy iterative processes (29, 32), which make them impractical to use on large scale clinical datasets.

Finally, we have demonstrated the true potential of *SynthSeg*⁺ on more than 14,000 uncurated clinical scans, where it accurately replicates volume trajectories observed on data of much higher quality. This result suggests that *SynthSeg*⁺ can be used to investigate other population effects on huge amounts of clinical data, which will considerably increase the statistical power of the current research studies. Moreover, *SynthSeg*⁺ also unlocks other potential applications, such as the analysis of scans acquired with the promising portable low-field MRI scanners, or the introduction of quantitative morphometry in the clinic for diagnosing and monitoring diseases.

We emphasize that *SynthSeg*⁺ can be run “out of the box” on brain MRI scans of any contrast and resolution, which has three main benefits. First, it greatly facilitates the use of our method since it eliminates the need for retraining and, thus, the associated requirements in terms of labeled data, deep-learning expertise, and hardware. Second, relying on a single model considerably improves the reproducibility of the results, since no hyperparameter tuning is required. And finally, it makes *SynthSeg*⁺ easier to disseminate, which is done here by distributing it with the publicly available package FreeSurfer.

Future work will first focus on surface placement to estimate cortical thickness, which is a powerful biomarker for the progression of many neuropsychiatric disorders. While this task may be challenging when analyzing low-resolution clinical scans (since the slice thickness can be locally larger than cortical folds), we believe that measurements can still be extracted in regions where the surface is orthogonal to the acquisition direction.

Then, we will seek to extend *SynthSeg*⁺ to analyze multimodal data. So far, we have dealt with multimodal acquisitions by processing all channels separately, but combining them into a common framework could improve accuracy. Finally, even though the aging clinical study has demonstrated the applicability of *SynthSeg*⁺ to large cohorts with high morphological variability, future work will aim to precisely quantify performance against various pathologies.

Overall, by enabling robust and reproducible analysis of heterogeneous clinical brain MRI scans, we believe that the present work will enable the development of clinical neuroimaging studies with sample sizes considerably higher than those found in research, which has the potential to revolutionize our understanding of the healthy and diseased human brain.

Materials and Methods

Training Datasets and Population Robustness. The proposed method is trained only on synthetic data (no real images) generated from a set of brain segmentation maps. Here, we use 1,020 maps obtained from 1 mm T1-weighted scans: 20 from the OASIS database (43), 500 from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) (7), and 500 from the Human Connectome Project (HCP) (44). These segmentations contain labels for 31 brain structures, obtained manually (OASIS) or with FreeSurfer (HCP and ADNI) (45). Moreover, we complement each map with 11 automated labels for extra-cerebral regions (46) and 68 FreeSurfer labels for cortex parcellation. We emphasize that, while HCP subjects are young and healthy, ADNI contains aging and diseased subjects, who frequently exhibit large atrophy and white matter lesions. Thus, using such a diverse population enables us to build robustness across a wide range of morphologies.

Brain MRI Test Datasets. Our experiments feature three datasets. The first one comprises 15,346 clinical scans from the PACS of Massachusetts General Hospital (see detailed information in *SI Appendix, Appendix 1*). Briefly, these scans are from 1,367 MRI sessions of distinct subjects with memory complaints, between 18 and 90 years of age: 749 males (age = 62.2 ± 15.2) and 618 females (age = 58.1 ± 17.2). Importantly, all scans are uncurated, and span a huge range of MR contrasts (T1-weighted, T2-weighted, FLAIR, diffusion MRI, etc.). Acquisitions are isotropic (11%) and anisotropic in axial (81%), coronal (4%), and sagittal (4%) orientations. The resolution of isotropic scans varies between

0.3 and 4.7 mm. For anisotropic scans, in-plane resolution ranges between 0.2 and 4.7 mm, while slice spacing varies between 0.8 and 10.5 mm. Ground truths for whole-brain segmentation, cortex parcellation, and ICV estimation were obtained for a subset of scans as follows. First, we isolated all sessions ($N = 62$) with 1 mm isotropic T1-weighted scans. These were then labeled using FreeSurfer (3), and the obtained segmentations were rigidly registered (47) to the other scans of the corresponding sessions. Note that we also obtained ICVs for all these scans by reporting the estimations given by FreeSurfer on the T1-weighted scans. Finally, we conducted a visual QC on the results and removed all scans where even a small segmentation error could be seen, either due to FreeSurfer or registration errors (138 scans) and/or poor image quality (94 cases with, e.g., insufficient coverage of the brain, wrong organ). In total, this provided us with ground truth segmentations for 520 scans, that we split between validation (20), and testing (500). All the other 14,752 scans were held-out for indirect evaluation.

The second dataset consists of 66 scans from three subdatasets: 20 T1-weighted scans from OASIS database (43); 18 subjects imaged twice with T2-weighted acquisitions and a sequence typically used in deep brain stimulation (DBS) (48); and 8 proton density scans (49). All scans are at 1 mm isotropic resolution and are available with manual or semiautomated segmentations for 31 brain regions (45). Labels for cortex parcellation were obtained by running FreeSurfer on the T1-weighted scans or on companion 1 mm T1-weighted acquisitions for the T2-weighted, DBS, and proton density scans.

The last dataset is another subset of 100 scans from ADNI (7), including 47 males and 53 females, aged 72.9 ± 7.6 y. Half of the subjects are healthy, while the others are diagnosed with Alzheimer's Disease. All subjects are imaged with two acquisitions: 1 mm isotropic T1-weighted scans and FLAIR scans at 5 mm axial resolution. Volumetric measurements for individual regions and ICVs are retrieved for all subjects by processing the T1-weighted scans with FreeSurfer.

Deep Hierarchical Segmentation. The proposed architecture relies on four CNN modules, which are designed to split the target segmentation task into easier intermediate operations. Specifically, a first segmenter S_1 is trained to produce initial segmentations of four coarse labels (cerebral white matter, cerebral gray matter, cerebrospinal fluid, and cerebellum). These labels group brain regions of similar tissue types and intensities and are thus easier to discriminate than individual structures. The output of S_1 is then fed to a denoising network D (28), which seeks to correct potential topological inconsistencies and larger segmentation mistakes that sometimes occur for scans with poor tissue contrast or low signal-to-noise ratio. Final whole-brain segmentations are obtained by a second segmenter S_2 , which takes as inputs the image and the corrected preliminary segmentations. As such, S_2 learns to segment the 31 individual target regions by using the coarse tissue segmentations as priors. We note that S_2 is given the opportunity to refine the boundaries produced by D , which are sometimes excessively smooth. Finally, the segmentations of S_2 are completed by passing them to a fourth network S_3 , which is trained to parcellate the left and right cortex labels from S_2 into 68 different substructures.

Training of the Segmenters. The three segmentation CNNs are trained separately with extremely diverse unimodal synthetic data created from anatomical segmentation maps by using a generative model inspired by Bayesian segmentation. Crucially, in order to train domain-agnostic networks, we adopt the domain randomization strategy introduced in our previous work (21). Specifically, we draw all the parameters governing the generative model at each minibatch from uniform priors of very large range. Hence, the segmenters are exposed to vastly changing examples in terms of shape, MR contrast, resolution, artifacts (bias field, noise), and even morphology (due to population variability in the training set). As a result, the segmenters are forced to learn contrast and resolution-agnostic features, which enables them to be applied to brain MRI scans of any domain, without requiring any retraining.

Briefly, this generative model requires only a set of 1 mm isotropic label maps as inputs, and synthesizes training examples as follows. At each minibatch, we randomly select one of the training maps and geometrically augment it with a random spatial transform (19). Next, a preliminary image is built by sampling

a randomized Gaussian mixture model conditionally on the deformed label map (50). The resulting image is then obtained by consecutively applying a random bias field, noise injection, intensity rescaling between 0 and 1, and a random voxel-wise exponentiation (19). In turn, low-resolution and PV effects are modeled with Gaussian blurring and subsampling at random low resolution. Finally, training pairs are obtained by defining the deformed label map as ground truth and resampling the low-resolution images back to the 1 mm isotropic grid, such that the downstream segmenters are trained to operate at high resolution (51).

Training of the Denoiser. State-of-the-art methods for denoising and topological correction rely mainly on supervised CNNs that are trained to recover ground truth segmentations from artificially corrupted versions of the same maps (18, 28, 52). However, the strategies used to corrupt the input segmentations are often handcrafted (random erosion and dilation, swapping of labels, etc.) and thus do not accurately capture the type of errors made by the segmentation method to correct. In order to train D with examples that are representative of S_1 errors, we instead degrade real images and feed them to the trained (and frozen) S_1 . D is then trained to map the outputs of S_1 back to their ground truth. During training, images are degraded on the fly with operations similar to the training of segmenters: spatial deformation, bias field, voxel-wise exponentiation, simulation of low resolution, and noise injection. However, the ranges of the parameters controlling these corruptions are made considerably wider than for the training of segmenters, in order to more frequently obtain erroneous segmentations from S_1 and thus to enrich the training data.

Automated QC Module. While the proposed architecture considerably improves robustness, it remains important to detect potential erroneous predictions, especially when segmenting clinical scans of varying quality. Hence, we introduce another module for automated failure detection. More precisely, we train a regressing network R to predict "performance scores" for 10 representative regions of interest (white matter, cortex, lateral ventricle, cerebellum, thalamus, hippocampus, amygdala, pallidum, putamen, brainstem), based solely on the segmentations produced by S_2 . Here, the performance scores aim to reflect Dice scores that would have been obtained if the input scans were available with associated ground truths (53). The segmentation of a region is then classified as failed if it obtains a predicted Dice score lower than 0.65 (a value chosen based on the validation set). In practice, R is trained with the same method as D , where we degrade real images, segment them with S_2 , and feed the obtained segmentations to R .

Network Architectures. All segmenters use the same architecture as *Synth-Seg* (50), which is based on a 3D UNet (14). Briefly, it comprises five levels, each consisting of two convolutions, a batch normalization, and either a max-pooling (contracting path), or upsampling operation (expanding path). Every convolution employs a $3 \times 3 \times 3$ kernel and an Exponential Linear Unit activation (54), except for the last layer, which uses a softmax. The first layer has 24 features, while this number is doubled after each max-pooling and halved after each upsampling. Finally, following the UNet architecture, we use skip connections across the contracting and expanding paths.

The denoiser uses a similar, but slightly lighter architecture: it only has one convolution per level and keeps a constant number of 16 features. Moreover, we suppress the skip connections between the top two levels to find a compromise between UNets, where top-level skip connections can potentially reintroduce erroneous features at late stages of the network; and auto-encoders, with excessive bottleneck-induced smoothness.

Finally, the regression network follows the same architecture as the encoder of the segmentation CNNs, except that it uses $5 \times 5 \times 5$ convolutions as in ref. 32, which greatly improved the results on the validation set. Regression scores are then retrieved by appending two more convolutions of 10 features (one for each QC region) and a global max-pooling.

Learning. The three segmenters and the denoiser are trained to minimize the average soft Dice loss (13). If Y_k represents the soft prediction for label $k \in \{1, \dots, K\}$, and T_k is its associated ground truth, this loss is given by

$$\text{Loss}_{\text{seg}} = 1 - \frac{1}{K} \sum_{k=1}^K \frac{2 \times \sum_{x,y,z} Y_k(x,y,z) T_k(x,y,z)}{\sum_{x,y,z} Y_k(x,y,z)^2 + T_k(x,y,z)^2} \quad [1]$$

The regressor is trained using a sum of square loss function. All networks are trained separately with the Adam optimizer (55) using a learning rate of 10^{-5} . We train each module until convergence, which approximately takes 300,000 steps for the segmenters (seven days on a Nvidia RTX6000) and 50,000 steps for the regressor (one day on the same GPU). All models are implemented in Keras (56) with a Tensorflow backend (57).

Inference. Test scans are automatically resampled to 1 mm isotropic resolution, and their intensities are normalized between 0 and 1. The trained model then predicts soft probabilistic segmentations for all target labels. Finally, hard segmentations are obtained by applying an argmax operation on these soft predictions. Overall, the whole inference pipeline takes between 12 and 16 s per scan on a RTX6000 GPU. We emphasize that when presented with multimodal data, SynthSeg+ segments each channel separately.

Individual Volumes and ICV Estimation. Volumes of individual brain regions are estimated by summing all the values of the corresponding soft predictions and multiplying the result by the volume of a voxel (1 mm^3). Note that summing soft probabilities rather than hard segmentations enables us to account for segmentation uncertainties and, to a certain extent, for partial voluming (12). In turn, ICVs are estimated for every subject by summing the predicted volumes of all structures, including the intracranial cerebro-spinal fluid.

Dice Scores. Parts of the evaluations use hard Dice scores, which measure the overlap between the same region across two hard segmentations. If X and Y are corresponding structures in two segmentations, their hard Dice score is given by

$$\text{Dice}(X, Y) = 2 \times \frac{|X \cap Y|}{|X| + |Y|}, \quad [2]$$

where $|\cdot|$ represents the cardinality of a set. Therefore, Dice scores vary between 0 (no overlap) and 1 (perfect matching).

1. J. Ashburner, K. Friston, Unified segmentation. *NeuroImage* **26**, 39–51 (2005).
2. M. Jenkinson, C. Beckmann, T. Behrens, M. Woolrich, S. Smith, FSL. *NeuroImage* **62**, 782–790 (2012).
3. B. Fischl, FreeSurfer. *NeuroImage* **62**, 774–781 (2012).
4. O. Oren, E. Kebebew, J. Ioannidis, Curbing unnecessary and wasted diagnostic imaging. *JAMA* **321**, 245–246 (2019).
5. D. Hibar *et al.*, Common genetic variants influence human subcortical brain structures. *Nature* **520**, 224–229 (2015).
6. F. Alfaro-Almagro *et al.*, Image processing and Quality Control for the first 10,000 brain imaging datasets from UK Biobank. *NeuroImage* **166**, 400–424 (2018).
7. C. R. Jack *et al.*, The Alzheimer's Disease Neuroimaging Initiative (ADNI): MRI methods. *J. Magn. Reson. Imaging: JMRI* **27**, 685–691 (2008).
8. Alzheimer's Disease Neuroimaging Initiative, ADNI participant demographic. <https://adni.loni.usc.edu/data-samples/adni-participant-demographic/> (Accessed 1 February 2023).
9. A. Fry *et al.*, Comparison of sociodemographic and health-related characteristics of UK biobank participants with those of the general population. *Am. J. Epidemiol.* **186**, 1026–1034 (2017).
10. O. Puonti, J. E. Iglesias, K. Van Leemput, Fast and sequence-adaptive whole-brain segmentation using parametric Bayesian modeling. *NeuroImage* **143**, 235–249 (2016).
11. H. Choi, D. Haynor, Y. Kim, Partial volume tissue classification of multichannel magnetic resonance images—a mixel model. *IEEE Trans. Med. Imaging* **10**, 395–407 (1991).
12. K. Van Leemput, F. Maes, D. Vandermeulen, P. Suetens, A unifying framework for partial volume segmentation of brain MR images. *IEEE Trans. Med. Imaging* **22**, 105–119 (2003).
13. F. Milletari, N. Navab, S. Ahmadi, "V-Net: Fully convolutional neural networks for volumetric medical image segmentation" in *International Conference on 3D Vision* (2016), pp. 565–571.
14. O. Ronneberger, P. Fischer, T. Brox, "U-Net: Convolutional networks for biomedical image segmentation" in *Medical Image Computing and Computer Assisted Intervention* (2015), pp. 234–241.
15. S. Pan, Q. Yang, A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* **22**, 45–59 (2010).
16. M. Ghafoorian *et al.*, "Transfer learning for domain adaptation in MRI: Application in brain lesion segmentation" in *Medical Image Computing and Computer Assisted Intervention* (2017), pp. 516–524.
17. Z. Akkus, A. Galimzianova, A. Hoogi, D. Rubin, B. Erickson, Deep learning for brain MRI segmentation: State of the art and future directions. *J. Digital Imaging* **30**, 449–459 (2017).
18. N. Karani, E. Erdil, K. Chaitanya, E. Konukoglu, Test-time adaptable neural networks for robust medical image segmentation. *Med. Image Anal.* **68**, 2–15 (2021).

Cohen's d. In this work, effect sizes in hippocampal volume between control and AD populations are measured with Cohen's d (38). If μ_C , s_C^2 and μ_{AD} , s_{AD}^2 designate the sample means and variances of two volume populations of size n_C and n_{AD} , where C stands for Controls, Cohen's d is given by

$$d = \frac{|\mu_C - \mu_{AD}|}{s}, \quad s = \sqrt{\frac{(n_C - 1)s_C^2 + (n_{AD} - 1)s_{AD}^2}{n_C + n_{AD} - 2}}. \quad [3]$$

Cohen's d below 0.2 are considered to be small, whereas values above 0.8 indicate large effect sizes (38).

Regression Model for Aging. Our aging model includes B-splines with 10 equally spaced knots for age, linear terms for slice spacing in each acquisition direction (i.e., sagittal, coronal, and axial) and a bias for gender. We then fit this model numerically by minimizing the sum of squares of the residuals with the L-BFGS-B method (58).

Data, Materials, and Software Availability. The code is available at <https://github.com/BBillot/SynthSeg>, and the trained models for SynthSeg+ and SynthSeg are distributed with the publicly available neuroimaging package FreeSurfer (3). The clinical dataset taken from the Massachusetts General Hospital is strictly private and cannot be shared. The remaining datasets used here are taken from databases that were made publicly available by their owners: ADNI (<https://adni.loni.usc.edu>), HCP (<https://www.humanconnectome.org>), OASIS (<https://www.oasis-brains.org>).

ACKNOWLEDGMENTS. This work is supported by the European Research Council (ERC Starting Grant 677697), the EPSRC UCL Centre for Doctoral Training in Medical Imaging (EP/L016478/1), the Department of Health's NIHR-funded Biomedical Research Centre UCLH, Alzheimer's Research UK (ARUK-IRG2019A-003), and the NIH (1R01AG070988, 1R15MH12319501).

Author affiliations: ^aCentre for Medical Image Computing, University College London, London WC1V 6LJ, UK; ^bDepartment of Neurology, Massachusetts General Hospital, Harvard Medical School, Boston, MA 02114; ^cMartinos Center for Biomedical Imaging, Massachusetts General Hospital, Harvard Medical School, Cambridge, MA 02129; and ^dComputer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02138

19. L. Zhang *et al.*, Generalizing deep learning for medical image segmentation to unseen domains via deep stacked transformation. *IEEE Trans. Med. Imaging* **39**, 2531–2540 (2020).
20. C. Chen, Q. Dou, H. Chen, J. Qin, P. A. Heng, Synergistic image and feature adaptation: Towards cross-modality domain adaptation for medical image segmentation. *AAAI Conf. Artif. Intell.* **33**, 65–72 (2019).
21. B. Billot *et al.*, SynthSeg: Domain Randomisation for Segmentation of Brain Scans of any Contrast and Resolution arXiv [Preprint] [cs] (2021). <http://arxiv.org/abs/2107.09559> (Accessed 4 January 2023).
22. J. Tobin *et al.*, "Domain randomization for transferring deep neural networks from simulation to the real world" in *IEEE/RSJ International Conference on Intelligent Robots and Systems* (2017), pp. 23–30.
23. F. Isensee, P. Jaeger, S. Kohl, J. Petersen, K. Maier-Hein, nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. *Nat. Methods* **18**, 203–211 (2021).
24. H. R. Roth *et al.*, An application of cascaded 3D fully convolutional networks for medical image segmentation. *Comput. Med. Imaging Graphics* **66**, 90–99 (2018).
25. M. Nosrati, G. Hamarneh, Incorporating prior knowledge in medical image segmentation: A survey arXiv [Preprint] [cs] (2016). <http://arxiv.org/abs/1607.01092> (Accessed 15 July 2016).
26. O. Oktay *et al.*, Anatomically constrained neural networks (ACNNs): Application to cardiac image enhancement and segmentation. *IEEE Trans. Med. Imaging* **37**, 384–395 (2018).
27. M. Sinclair *et al.*, Atlas-ISTN: Joint segmentation, registration and atlas construction with image- and spatial transformer networks. *Med. Image Anal.* **78**, 102–123 (2022).
28. A. Larrazabal, C. Martinez, B. Glocker, E. Ferrante, Post-DAE. *IEEE Trans. Med. Imaging* **39**, 3813–3820 (2020).
29. V. Valindria *et al.*, Reverse classification accuracy: Predicting segmentation performance in the absence of ground truth. *IEEE Trans. Med. Imaging* **36**, 1597–1606 (2017).
30. T. Kohlberger, V. Singh, C. Alvino, C. Bahlmann, L. Grady, "Evaluating segmentation error without ground truth" in *Medical Image Computing and Computer Assisted Intervention* (2019), pp. 528–536.
31. F. Liu, Y. Xia, D. Yang, A. Yuille, D. Xu, "An alarm system for segmentation algorithm based on shape model" in *ICCV* (2019), pp. 10652–10661.
32. S. Wang *et al.*, "Deep generative model-based quality control for cardiac MRI segmentation" in *Medical Image Computing and Computer Assisted Intervention* (2020), pp. 88–97.
33. B. Billot, C. Magdamo, S. E. Arnold, S. Das, J. E. Iglesias, "Robust segmentation of brain MRI in the wild with hierarchical CNNs and no retraining" in *Medical Image Computing and Computer Assisted Intervention* (2022), pp. 538–548.

34. E. DeLong, D. DeLong, D. Clarke-Pearson, Comparing the areas under two or more correlated receiver operating characteristic curves: A nonparametric approach. *Biometrics* **44**, 837-845 (1988).
35. A. P. Bradley, The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognit.* **30**, 1145-1159 (1997).
36. I. Malone *et al.*, Accurate automatic estimation of total intracranial volume: A nuisance variable with less nuisance. *NeuroImage* **104**, 366-372 (2015).
37. K. Gosche, J. Mortimer, C. Smith, W. Markesbery, D. Snowdon, Hippocampal volume as an index of Alzheimer neuropathology. *Neurology* **58**, 1476-1482 (2002).
38. J. Cohen, *Statistical Power Analysis for the Behavioural Sciences* (Routledge Academic, 1988).
39. R. Bethlehem, J. Seidlitz, S. White, J. Vogel, K. Anderson, Brain charts for the human lifespan. *Nature* **604**, 525-533 (2022).
40. P. Coupé, G. Catheline, E. Lanuza, J. Manjón, Towards a unified analysis of brain maturation and aging across the entire lifespan: A MRI analysis. *Hum. Brain Mapp.* **38**, 5501-5518 (2017).
41. D. Dima *et al.*, Subcortical volumes across the lifespan: Data from 18,605 healthy individuals aged 3-90 years. *Hum. Brain Mapp.* **43**, 452-469 (2022).
42. L. Henschel *et al.*, Fastsurfer - a fast and accurate deep learning based neuroimaging pipeline. *NeuroImage* **219**, 117-136 (2020).
43. D. Marcus *et al.*, Open access series of imaging studies (OASIS): Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *J. Cognit. Neurosci.* **19**, 1498-1507 (2007).
44. D. Van Essen *et al.*, The Human Connectome Project: A data acquisition perspective. *NeuroImage* **62**, 2222-2231 (2012).
45. B. Fischl *et al.*, Whole brain segmentation: Automated labeling of neuroanatomical structures in the human brain. *Neuron* **33**, 41-55 (2002).
46. O. Puonti *et al.*, Accurate and robust whole-head segmentation from magnetic resonance images for individualized head modeling. *NeuroImage* **219**, 317-332 (2020).
47. M. Modat *et al.*, Fast free-form deformation using graphics processing units. *Comput. Methods Progr. Biomed.* **98**, 278-284 (2010).
48. J. E. Iglesias *et al.*, A probabilistic atlas of the human thalamic nuclei combining ex vivo MRI and histology. *NeuroImage* **183**, 314-326 (2018).
49. B. Fischl *et al.*, Sequence-independent segmentation of magnetic resonance images. *NeuroImage* **23**, 69-84 (2004).
50. B. Billot *et al.*, "A learning strategy for contrast-agnostic MRI segmentation" in *Medical Imaging with Deep Learning* (2020), pp. 75-93.
51. B. Billot, E. Robinson, A. Dalca, J. E. Iglesias, "Partial Volume segmentation of brain MRI scans of any resolution and contrast" in *Medical Image Computing and Computer Assisted Intervention* (2020), pp. 177-187.
52. M. Khan, M. Gajendran, Y. Lee, M. Khan, Deep neural architectures for medical image semantic segmentation: Review. *IEEE Access* **9**, 83002-83024 (2021).
53. E. Hann, L. Biasioli, Q. Zhang, S. Neubauer, S. Piechnik, "Quality control-driven image segmentation: Towards reliable automatic image analysis in large-scale cardiovascular magnetic resonance aortic cine imaging" in *Medical Image Computing and Computer Assisted Intervention* (2019), pp. 750-758.
54. D. A. Clevert, T. Unterthiner, S. Hochreiter, Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs). *arXiv [Preprint] [cs]* (2016). <http://arxiv.org/abs/1511.07289> (Accessed 22 February 2016).
55. D. Kingma, J. Ba, Adam: A Method for Stochastic Optimization. *arXiv [Preprint] [cs]* (2017). <http://arxiv.org/abs/1412.6980> (Accessed 30 January 2017).
56. F. Chollet, Keras, Version 2.3.1. <https://keras.io> (Accessed 1 February 2023).
57. M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, "Tensorflow: A system for large-scale machine learning" in *Symposium on Operating Systems Design and Implementation* (2016), pp. 265-283.
58. R. Byrd, P. Lu, J. Nocedal, C. Zhu, A limited memory algorithm for bound constrained optimization. *J. Sci. Comput.* **16**, 1190-1208 (1995).