



## Research article

# Hippocampus segmentation on epilepsy and Alzheimer's disease studies with multiple convolutional neural networks <sup>☆</sup>



Diedre Carmo <sup>a,\*</sup>, Bruna Silva <sup>b</sup>, Alzheimer's Disease Neuroimaging Initiative <sup>1</sup>, Clarissa Yasuda <sup>b</sup>, Letícia Rittner <sup>a</sup>, Roberto Lotufo <sup>a</sup>

<sup>a</sup> School of Electrical and Computer Engineering, UNICAMP, Campinas, São Paulo, Brazil

<sup>b</sup> Faculty of Medical Sciences, UNICAMP, Campinas, São Paulo, Brazil

## ARTICLE INFO

## Keywords:

Deep learning  
Hippocampus segmentation  
Convolutional neural networks  
Alzheimer's disease  
Epilepsy

## ABSTRACT

**Background:** Hippocampus segmentation on magnetic resonance imaging is of key importance for the diagnosis, treatment decision and investigation of neuropsychiatric disorders. Automatic segmentation is an active research field, with many recent models using deep learning. Most current state-of-the-art hippocampus segmentation methods train their methods on healthy or Alzheimer's disease patients from public datasets. This raises the question whether these methods are capable of recognizing the hippocampus on a different domain, that of epilepsy patients with hippocampus resection.

**New Method:** In this paper we present a state-of-the-art, open source, ready-to-use, deep learning based hippocampus segmentation method. It uses an extended 2D multi-orientation approach, with automatic pre-processing and orientation alignment. The methodology was developed and validated using HarP, a public Alzheimer's disease hippocampus segmentation dataset.

**Results and Comparisons:** We test this methodology alongside other recent deep learning methods, in two domains: The HarP test set and an in-house epilepsy dataset, containing hippocampus resections, named HCUnicamp. We show that our method, while trained only in HarP, surpasses others from the literature in both the HarP test set and HCUnicamp in Dice. Additionally, Results from training and testing in HCUnicamp volumes are also reported separately, alongside comparisons between training and testing in epilepsy and Alzheimer's data and vice versa.

**Conclusion:** Although current state-of-the-art methods, including our own, achieve upwards of 0.9 Dice in HarP, all tested methods, including our own, produced false positives in HCUnicamp resection regions, showing that there is still room for improvement for hippocampus segmentation methods when resection is involved.

## 1. Introduction

The hippocampus is a small, medial, subcortical brain structure related to long and short term memory [1]. The hippocampus can be affected in shape and volume by different pathologies, such as the neurodegeneration associated to Alzheimer's disease [2], or surgical intervention to treat temporal lobe epilepsy [3]. Hippocampal segmentation from magnetic resonance imaging (MRI) is of great importance for research of neuropsychiatric disorders and can also be used in the preoperative investigation of pharmaco-resistant temporal lobe epilepsy [4].

The medical research of these disorders usually involves manual segmentation of the hippocampus, requiring time and expertise in the field. The high-cost associated to manual segmentation has stimulated the search for effective automatic segmentation methods. Some of those methods, such as FreeSurfer [5], are already used as a starting point for a manual finer segmentation later [6].

While conducting research on epilepsy and methods for hippocampus segmentation, two things raised our attention. Firstly, the use of deep learning and Convolutional Neural Networks (CNN) is in the spotlight, with most of the recent hippocampus segmentation methods

<sup>☆</sup> Source code: <https://github.com/MICLab-Unicamp/e2dhipseg>.

\* Corresponding author.

E-mail address: [diedre@dca.fee.unicamp.br](mailto:diedre@dca.fee.unicamp.br) (D. Carmo).

<sup>1</sup> Part of the data used in preparation of this article were obtained from the Alzheimer's Disease Neuroimaging Initiative (ADNI) database ([adni.loni.usc.edu](http://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](http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf).

<https://doi.org/10.1016/j.heliyon.2021.e06226>

Received 25 September 2020; Received in revised form 6 December 2020; Accepted 3 February 2021

featuring CNNs. Secondly, many of these methods rely on publicly available datasets for training and evaluating and therefore have access only to healthy scans, or patients with Alzheimer's disease. This raises the concern that automated methods might only be prepared to deal with features present in the public Alzheimer's and healthy subjects datasets, such as ADNI and the Multi Atlas Labeling Challenge (MALC).

Considering these facts, we present an improved version of our own deep learning based hippocampus segmentation method [25], compared with other recent methods [7, 8, 9]. We use the public Alzheimer's HarP dataset for training and initial testing comparisons with other methods. As an additional test dataset, an in-house epilepsy dataset named HCUnicamp is used. It contains scans from patients with epilepsy (pre and post surgical removal of hippocampus), with different patterns of atrophy compared to that observed both in the Alzheimer's data and healthy subjects. It is important to note that HCUnicamp is not involved in our method's training or methodological choices, to allow for fair comparisons with other methods. Without comparing to other methods, we also report results of involving HCUnicamp epilepsy volumes in training.

### 1.1. Contributions

In summary, the main contributions of this paper are as follows:

- A readily available hippocampus segmentation methodology under the MIT license, consisting of an ensemble of 2D CNNs coupled with traditional 3D post processing, achieving state of the art performance in HarP public data, and using recent advancements from the deep learning literature.
- An evaluation of recent hippocampus segmentation methods in our epilepsy test dataset, HCUnicamp, that includes post-operative images of patients without one of the hippocampi. In this evaluation, our method is only trained in public HarP volumes, therefore our methodology has no bias related to this task. We show that our method is also superior in this domain, although no method was able to achieve more than 0.8 Dice in this dataset, according to our manual annotations. As far as we know, that has not been explored before with recent Deep Learning methods.
- A final experiment includes epilepsy HCUnicamp volumes in training, without changing the methodology and with no comparisons to other methods, which resulted in better performance on epilepsy cases. The effects of mixing data from both datasets in training are explored.

This paper is organized as follows: Section 2 presents a literature review of recent deep learning based hippocampus segmentation methods. Section 3 introduces more details to the two datasets involved in this research. A detailed description of our hippocampus segmentation methodology is in Section 4. Section 5 has experimental results from our methodology development, qualitative and quantitative comparisons with other methods in HarP and HCUnicamp, and results of involving HCUnicamp volumes in training. Sections 6 and 7 have, respectively, extended discussion of those results and conclusion. More details to the training and hyperparameter optimization process are in the appendix.

## 2. Hippocampus segmentation with deep learning

Before the rise of deep learning methods in medical imaging segmentation, most hippocampus segmentation methods used some form of optimization of registration and deformation to atlas(es) [5, 10, 11, 12, 13, 14]. Even today, medical research uses results from FreeSurfer [5], a high impact multiple brain structures segmentation work, available as a software suite. Those atlas-based methods can produce high quality segmentations, taking, however, around 8 hours in a single volume. Lately, a more time efficient approach appeared in the literature,

namely the use of such atlases as training volumes for CNNs. Deep learning methods can achieve similar overlap metrics while predicting results in a matter of seconds per volume [7, 8, 15, 16, 17, 18, 19].

Recent literature on hippocampus segmentation with deep learning is exploring different architectures, loss functions and overall methodologies for the task. One approach that seems to be common to most of the studies involves the combination of 2D or 3D CNNs, and patches as inputs in the training phase. Note that some works focus on hippocampus segmentation, while some attempt segmentation of multiple neuroanatomy. Following, a brief summary of each of those works.

Chen et al. [15] reports 0.9 Dice [20] in 10-fold 110 ADNI [2] volumes with a novel CNN input idea. Instead of using only the triplanes as patches, it also cuts the volume in six more diagonal orientations. This results in 9 planes, that are fed to 9 small modified U-Net [21] CNNs. The ensemble of these U-Nets constructs the final result.

Xie et al. [16] trains a voxel-wise classification method using triplanar patches crossing the target voxel. They merge features from all patches into a Deep Neural Network with a fully connected classifier alongside standard use of ReLU activations and softmax [22]. The training patches come only from the approximate central area the hippocampus usually is, balancing labels for 1:1 foreground and background target voxels. Voxel classification methods tend to be faster than multi-atlas methods, but still slower than Fully Convolutional Neural Networks.

DeepNat from Wachinger et al. [17] achieves segmentation of 25 structures with a 3D CNN architecture. With a hierarchical approach, a 3D CNN separates foreground from background and another 3D CNN segments the 25 sub-cortical structures on the foreground. Alongside a proposal of a novel parametrization method replacing coordinate augmentation, DeepNat uses 3D Conditional Random Fields as post-processing. The architecture is a voxel-wise classification, taking into account the classification of neighbor voxels. DeepNat mainly focuses on the MALC dataset, achieving around 0.86 Dice in hippocampus segmentation.

Thyreau et al. [8]'s model, named Hippodeep, uses CNNs trained in a region of interest (ROI). However, where we apply one CNN for each plane of view, Thyreau et al. uses a single CNN, starting with a planar analysis followed by layers of 3D convolutions and shortcut connections. This study used more than 2000 patients, augmented to around 10000 volumes with augmentation. Initially the model is trained with FreeSurfer segmentations, and later fine-tuned using volumes which the author had access to manual segmentations, the gold standard. Thyreau's method requires MNI152 registration of input data, which adds around a minute of computation time, but the model is generally faster than multi-atlas or voxel-wise classification, achieving generalization in different datasets, as verified by Nogovitsyn et al. [23].

QuickNat from Roy et al. [7] achieves faster segmentations than DeepNat by using a multiple CNN approach instead of voxel-wise classification. Its methodology follows a consensus of multiple 2D U-Net like architectures specialized in each slice orientation. The use of FreeSurfer [5] masks over hundreds of public data to generate silver standard annotations allows for much more data than usually available for medical imaging. Later, after the network already knows to localize the structures, it is finetuned to more precise gold standard labels. Inputs for this method need to conform to the FreeSurfer format.

Ataloglou et al. [18] recently displayed another case of fusion of multiple CNN outputs, specialized into axial, coronal and sagittal orientations, into a final hippocampus segmentation. They used U-Net like CNNs specialized in each orientation, followed by error correction CNNs, and a final average fusion of the results. They went against a common approach in training U-Nets of using patches during data augmentation, instead using cropped slices. This raises concerns about overfitting to the used dataset, HarP [24], supported by the need of finetuning to generalize to a different dataset.

Dinsdale et al. [19] mixes knowledge from multi-atlas works with deep learning, by using a 3D U-Net CNN to predict a deformation field

from an initial binary sphere to the segmentation of the hippocampus, achieving around 0.86 DICE on HarP. Interestingly, trying an auxiliary classification task did not improve segmentation results.

It is known that deep learning approaches require a relatively large amount of varied training data. Commonly used forms of increasing the quantity of data in the literature include using 2D CNNs over regions (patches) of slices, with some form of patch selection strategy. The Fully Convolutional Neural Network (FCNN) U-Net [21] architecture has shown potential to learn from relatively small amounts of data with their decoding, encoding and concatenation schemes, even working when used with 3D convolutions directly in a 3D volume [9].

Looking at these recent works, one can confirm the segmentation potential of the U-Net architecture, including the idea of an ensemble of 2D U-Nets instead of using a single 3D one, as we [25], some simultaneous recent work [7, 18], or even works in other segmentation problems [26] presented. In this paper, some of those methods were reproduced for comparison purposes in our in-house dataset, namely [7, 8], including a 3D UNet architecture test from [9].

As far as we know, there is no study applying recent deep learning methods trained on public data, such as HarP and MALC, to MRI scans of epilepsy including hippocampus resection cases. We also include, separately, an attempt to train on such data.

### 3. Data

This study uses mainly two different datasets: one collected locally for an epilepsy study, named HCUnicamp; and one public from the ADNI Alzheimer's study, HarP. HarP is commonly used in the literature as a hippocampus segmentation benchmark. The main difference between the datasets is, the lack of one of the hippocampi in 70% of the scans from HCUnicamp, as these patients underwent surgical removal (Fig. 1).

Although our method needs input data to be in the MNI152 [27] orientation, data from those datasets are in native space and are not registered. We provide an automatic orientation correction by rigid registration as an option when predicting in external volumes, to avoid orientation mismatch problems.


#### 3.1. HarP

This methodology was developed with training and validation on HarP [24], a widely used benchmark dataset in the hippocampus segmentation literature. HarP uses data from the Alzheimer's disease Neuroimaging Initiative (ADNI) database ([adni.loni.usc.edu](http://adni.loni.usc.edu)). The ADNI was launched in 2003 as a public-private partnership, led by Principal Investigator Michael W. Weiner, MD. The primary goal of ADNI has been to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer's disease (AD).

The full HarP release contains 135 T1-weighted MRI volumes. Alzheimer's disease classes are balanced with equal occurrence of control normal (CN), mild cognitive impairment (MCI) and Alzheimer's disease (AD) cases [2]. Volumes were minmax intensity normalized between 0 and 1, and no volumes were removed. Training with stratified hold-out was performed with 80% training, 10% validation and 10% testing, while k-Folds, when used, consisted of 5 folds, with no overlap on the test sets.

#### 3.2. HCUnicamp

HCUnicamp was collected inhouse, by personnel from the Brazilian Institute of Neuroscience and Neurotechnology (BRAINN) at UNICAMP's *Hospital de Clínicas*. This dataset contains 190 T1-weighted 3T MRI



**Fig. 1.** (a) 3D rendering of the manual annotation (in green) of one of the HarP dataset volumes. In (b), a coronal center crop slice of the average hippocampus mask for all volumes in HarP (green) and HCUnicamp (red), shows different head alignment. Zero corresponds to the center. (c) Sagittal, (d) Coronal and (e) Axial HCUnicamp slices from a post-operative scan with annotations in red.

acquisitions, in native space. 58 are controls and 132 are epilepsy patients. From those epilepsy images, 70% had one of the hippocampi surgically removed, resulting in a very different shape and texture than what is commonly seen in public datasets (Fig. 1). More details about the surgical procedure can be found in [3, 4]. All volumes have manual annotations of the hippocampus, performed by one rater. The voxel intensity is minmax normalized, between 0 and 1, per volume. This data acquisition and use was approved by an Ethics and Research Committee (CEP/Conep, number 3435027).

A comparison between the datasets can be seen in Fig. 1. The difference in mean mask position due to the inclusion of neck in HCUnicamp is notable, alongside with the lower presence of left hippocampus labels due to surgical intervention for epilepsy (Fig. 1b).

To investigate the performance of different methods in terms of dealing with the absence of hippocampus and unusual textures, we used the whole HCUnicamp dataset (considered a different domain) as a final test set. Our methodology was only tested in this dataset at the end, alongside other methods. Results on HCUnicamp were not taken into consideration for our method's methodological choices, to allow for fair comparisons with other methods, treating this data as a true final test set.

A final additional experiment attempts to learn from the epilepsy data, dividing HCUnicamp in a balanced hold-out of 70% training, 10% validation and 20% testing. These subsets are called HCU-Train, HCU-Validation and HCU-Test for clarity.

### 4. Segmentation methodology

In this section, the general methodology (Fig. 2) for our hippocampus segmentation method is detailed. Three orientation specialized 2D U-Net CNNs are utilized, inspired by Lucena et al.'s work [26]. The activations from the CNNs are merged into an activation consensus. Each network's activations for a given input volume are built slice by slice. The three activation volumes are averaged into a consensus volume, which is post-processed into the final segmentation mask.

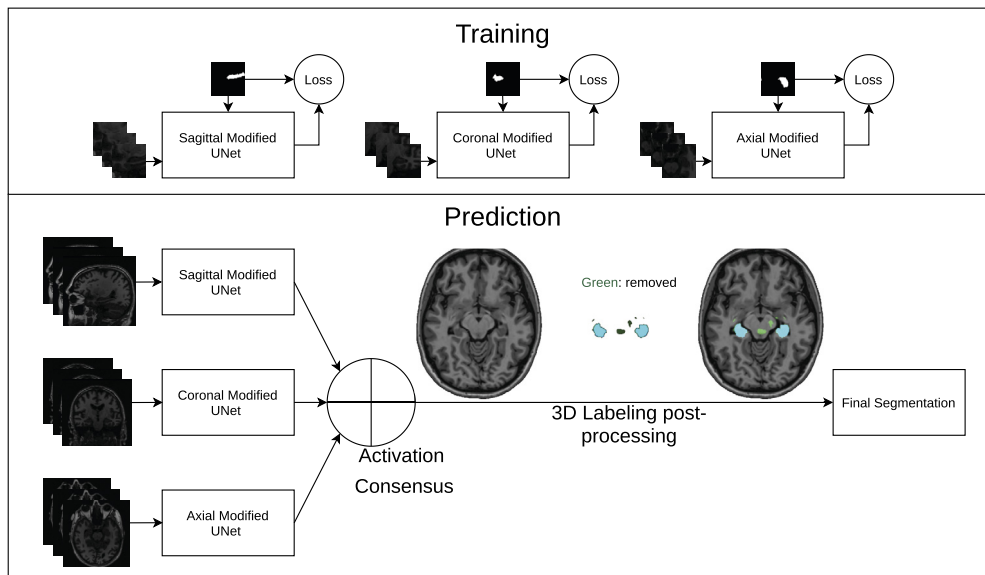


Fig. 2. The final segmentation volume is generated by taking into account activations from three FCNNs specialized on each 2D orientation. Neighboring slices are taken into account in a multi-channel approach. Full slices are used in prediction time, but training uses patches.

The basic structure of our networks is inspired by the U-Net FCNN architecture [21]. However, some modifications based on other successful works were applied to the architecture (Fig. 3). Those modifications include: instead of one single 2D patch as input, two neighbor patches are concatenated leaving the patch corresponding to the target mask in the center [28]. Residual connections based on ResNet [29] between the input and output of the double convolutional block were added, as 1x1 2D convolutions to account for different number of channels. Batch normalization was added to each convolution inside the convolutional block, to accelerate convergence and facilitate learning [30]. Also, all convolutions use padding to keep dimensions and have no bias. This works uses VGG11 [31] weights in the encoder part of the U-Net architecture, as in [32].

During prediction time, slices for each network are extracted with a center crop. When building the consensus activation volume, the resulting activation is padded back to the original size. For training, this method uses patches. Patches are randomly selected in runtime. Patches can achieve many possible sizes, as long as it accommodates the number of spatial resolution reductions present in the network.

A pre-defined percentage of the patches is selected from a random point of the brain, allowing for learning of what structures are not the hippocampus. Those are called negative patches. On the other hand, positive patches are always centered on a random point of the hippocampus border. In a similar approach to Pereira et al. [28]’s Extended 2D, adjacent patches (slices on evaluation) are included in the network’s input as additional channels (Fig. 2). The intention is for the 2D network to take into consideration volumetric information adjacent to the region of interest, hence the name for the method, Extended 2D Consensus Hippocampus Segmentation (E2DHipseg). This approach is inspired by how physicians compare neighbor slices in multiview visualization when deciding if a voxel is part of the analyzed structure or not.

Data augmentation is used to improve our dataset variance and avoid overfitting. All augmentations perform a random small runtime modification to the data. Random augmentations include intensity modification ([-0.05, 0.05]), rotation and scale ([-10, 10]) and gaussian noise with 0 mean and 0.0002 variance.

#### 4.1. Loss function

Dice [20] is an overlap metric widely used in the evaluation of segmentation applications. Performance in this paper is mainly evaluated

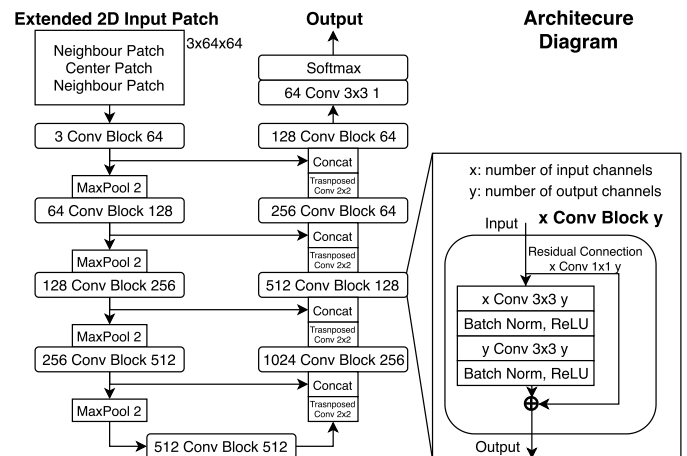


Fig. 3. Final architecture of each modified U-Net in Fig. 2. Of note in comparison to the original U-Net is the use of BatchNorm, residual connections in each convolutional block, the 3 channel neighbor patches input and the sigmoid output limitation. Padding is also used after convolutions.

with Dice, by comparisons with the manual gold standard. Dice can be defined as:

$$\frac{2 \sum_i^N p_i g_i}{\sum_i^N p_i^2 + \sum_i^N g_i^2} \quad (1)$$

where the sums run over the N voxels, of the predicted binary segmentation volume  $p_i \in P$  and the ground truth binary volume  $g_i \in G$ . For conversion from a metric to a loss function, one can simply optimize  $1 - Dice$ , therefore optimizing a segmentation overlap metric. This is referred here as Dice Loss.

To take into account background information, a Softmax of two-channels representing background and foreground can be used as an output. In this case, Generalized Dice Loss (GDL) [20] and Boundary Loss, a recent proposal of augmentation to GDL from Kervadec et al. [33] were considered as loss options.

Generalized Dice Loss weights the loss value by the presence of a given label in the target, giving more importance to less present labels. This solves the class imbalance problem that would emerge when using Dice Loss while including background as a class.



Boundary Loss takes into consideration alongside the “regional” loss (e.g. GDL), the distance between boundaries of the prediction and target, which does not give any weight to the area of the segmentation. Kervadec’s work suggests that a loss functions that takes into account boundary distance information can improve results, specially for unbalanced datasets. However, one needs to balance the contribution of both components with a weight, defined as  $\alpha$  in the following Boundary Loss (B) equation:

$$B(p, g) = \alpha G(p, g) + (1 - \alpha) S(p, g) \quad (2)$$

where G is GDL, regional component of the loss function, and S is the surface component, that operates on surface distances. The weight factor  $\alpha$  changes from epoch to epoch. The weight given to the regional loss is shifted to the surface loss, with  $\alpha$  varying from 1 in the first epoch to 0 in the last epoch. We followed the original implementation in [33].

#### 4.2. Consensus and post-processing

The consensus depicted in Fig. 2 consists of taking the average from the activations of all three CNNs. A more advanced approach of using a 4th, 3D, U-Net as the consensus generator was also attempted.

After construction of the consensus of activations, a threshold is needed to binarize the segmentation. We noticed that sometimes, small structures of the brain similar to the hippocampus could be classified as false positives. To remove those false positives, a 3D labeling implementation from [34] was used, with subsequent removal of small non-connected volumes, keeping the 2 largest volumes, or 1 if a second volume is not present (Fig. 2). This post processing is performed after the average consensus of all networks and threshold application.

## 5. Experiments and results

This section presents quantitative and qualitative comparisons with other methods in HarP and HCUnicamp. The appendix showcases more detailed experiments on the segmentation methodology, displaying differences in Dice in the HarP test set, resulting from our methodological choices.

### 5.1. Quantitative results

In this section, we report quantitative results of our method and others from the literature in both HarP and HCUnicamp. For comparison’s sake, we also trained an off-the-shelf 3D U-Net architecture, from Isensee et al. [9], originally a Brain Tumor segmentation work. Isensee’s architecture was trained with ADAM and HarP 3D center crops as input.

For the evaluation with the QuickNat [7] method, volumes and targets needed to be conformed to its required format, causing interpolation. As far as we know, the method does not have a way to return its predictions on the volume’s original space. DICE was calculated with the masks on the conformed space. Note that QuickNat performs segmentation of multiple brain structures.

#### 5.1.1. HarP

The best hold-out mean Dice is 0.9133. In regards to specific Alzheimer’s classes in the test set, our method achieves 0.9094 Dice for CN, 0.9378 for MCI and 0.9359 for AD cases. When using a hold-out approach in a relatively small dataset such as HarP, the model can be overfitted to better results in that specific test set. With that in mind, we also report results with cross validation. 5-fold training is used, applied to all three network’s training. With 5-fold our model achieved  $0.90 \pm 0.01$  Dice. Results reported by other works are present in Table 1. Our methodology has similar performance to what is reported by Ataloglou et al.’s recent, simultaneous work [18]. Interestingly, the initial methodology of both methods is similar, in the use of multiple 2D CNNs.

**Table 1.** Reported testing results for HarP. This work is named E2DHipseg. Results with \* were calculated following a 5-fold cross validation.

Deep learning methods	HarP (DICE)
3D U-Net - Isensee et al. [9] (2017)	0.86
Hippodeep - Thyreau et al. [8] (2018)	0.85
QuickNat - Roy et al. [7] (2019)	0.80
Ataloglou et al. [18] (2019)	0.90*
<b>E2DHipseg (this work)</b>	<b>0.90*</b>
Atlas-based methods	
FreeSurfer v6.0 [5] (2012)	0.70
Chincardini et al. [13] (2016)	0.85
Platero et al. [14] (2017)	0.85

#### 5.1.2. HCUnicamp

As described previously, the HCUnicamp dataset has lack of one of the hippocampi in many of its scans (Fig. 1), and it was used to examine the generalization capability of these methods. Table 2 has mean and standard deviation Dice for all HCUnicamp volumes, using both masks, or only one the left or right mask, with multiple methods. “with Aug.” refers to the use of augmentations in training. We also report Precision and Recall, per voxel classification, where positives are hippocampus voxels and negatives are non hippocampus voxels. Precision is defined by  $TP/(TP + FP)$  and Recall is defined by  $TP/(TP + FN)$ , where TP is true positives, FP are false positives and FN are false negatives. All tests were run locally. Unfortunately, we were not able to reproduce Ataloglou et al.’s method for local testing.

Our method performed better than other recent methods on the literature in the HCUnicamp dataset, even though HCUnicamp is not involved on our methodology development. However, no method was able to achieve more than 0.8 mean Dice in epilepsy patients. The high number of false positives due to hippocampus removal is notable by the low left and right DICE, and low precision. The impact of additional augmentations was not statistically significant in the epilepsy domain.

Our method takes around 15 seconds on a mid-range GPU and 3 minutes on a consumer CPU to run, per volume. All the code used on its development is available in [github.com/MICLab-Unicamp/e2dhipseg](https://github.com/MICLab-Unicamp/e2dhipseg), with instructions for how to run it in an input volume, under MIT license. A free executable version for medical research use, without environment dependencies, is available on the repository. To avoid problems with different head orientations, there is an option to use MNI152 registration when predicting in a given volume. Even when performing registration, the output mask will be in the input volume’s space, using the inverse transform. In regards to pre-processing requirements, our method requires only for the volume to be a 3D MRI in the correct orientation. The automatic MNI152 registration option solves this problem, in a similar way to Hippodeep. A GPU is recommended for faster prediction but not necessary.

#### 5.2. Adaptation to HCUnicamp

Additional experiments were performed now involving HCUnicamp data in training, to try and learn to recognize the resection. The experiments involved making a hold-out separation of HCUnicamp. In the previous experiment, all volumes were involved in the testing and not used for training of any method. In this one, hold-out with 70% training, 10% validation and 20% testing is performed with balance between control and patients, to allow for training. Note that these results are not comparable with other method’s results or even our own results present in Table 2, since the dataset is different and we are now training on part of HCUnicamp. To avoid confusion, the hold-out subsets will be referred to as HCU-Train and HCU-Test. Experiments were also performed including only control volumes or only patient volumes, with the same hold-out approach (Table 3). Results improve when training on HCUnicamp volumes, but the high standard deviation still shows that the method is failing to recognize resections.

**Table 2.** Locally executed testing results for HCUnicamp. All 190 volumes from the dataset are included, and no model saw it on training. The 3D U-Net here is using the same weights from table 1. QuickNat performs whole brain multitask segmentation, not only hippocampus.

HCUnicamp (Controls)					
Method	Both (Dice)	Left (Dice)	Right (Dice)	Precision	Recall
3D U-Net - Isensee et al. [9] (2017)	0.80 ± 0.04	0.81 ± 0.04	0.78 ± 0.04	0.76 ± 0.10	0.85 ± 0.06
Hippodeep - Thyreau et al. [8] (2018)	0.80 ± 0.05	0.81 ± 0.05	0.80 ± 0.05	0.72 ± 0.10	0.92 ± 0.04
QuickNat - Roy et al. [7] (2019)	0.80 ± 0.05	0.80 ± 0.05	0.79 ± 0.05	0.71 ± 0.11	<b>0.92 ± 0.04</b>
<b>E2DHipseg without Aug.</b>	0.82 ± 0.03	0.83 ± 0.03	0.82 ± 0.03	<b>0.78 ± 0.10</b>	0.88 ± 0.06
<b>E2DHipseg with Aug.</b>	<b>0.82 ± 0.03</b>	<b>0.83 ± 0.03</b>	<b>0.82 ± 0.04</b>	0.78 ± 0.10	0.89 ± 0.06
HCUnicamp (Patients)					
3D U-Net - Isensee et al. [9] (2017)	0.74 ± 0.08	0.48 ± 0.39	0.56 ± 0.36	0.66 ± 0.12	0.87 ± 0.07
Hippodeep - Thyreau et al. [8] (2018)	0.74 ± 0.08	0.48 ± 0.39	0.57 ± 0.37	0.63 ± 0.12	0.91 ± 0.06
QuickNat - Roy et al. [7] (2019)	0.71 ± 0.08	0.47 ± 0.38	0.56 ± 0.36	0.59 ± 0.12	<b>0.92 ± 0.06</b>
<b>E2DHipseg without Aug.</b>	<b>0.77 ± 0.07</b>	0.49 ± 0.40	0.58 ± 0.37	<b>0.69 ± 0.11</b>	0.88 ± 0.07
<b>E2DHipseg with Aug.</b>	0.76 ± 0.07	<b>0.50 ± 0.40</b>	<b>0.58 ± 0.37</b>	0.68 ± 0.11	0.89 ± 0.07

**Table 3.** E2DHipseg with networks trained in HCU-Train. Training is performed in all volumes, only patients or only controls, and testing is done in HCU-Test, also selecting for patients, controls, or all.

Trained on	Both (Dice)	Left (Dice)	Right (Dice)
Patients	0.84 ± 0.04	0.60 ± 0.41	0.56 ± 0.42
Controls and Patients	0.86 ± 0.05	0.71 ± 0.36	0.74 ± 0.34
Controls	0.90 ± 0.01	0.89 ± 0.02	0.90 ± 0.01

**Table 4.** This table compares training in one dataset and testing in the other. Better results are achieved when involving both domains in training.

Trained on	Tested on	Both (Dice)	Left (Dice)	Right (Dice)
Harp-Train	HCU-Test	0.79 ± 0.07	0.65 ± 0.33	0.68 ± 0.31
HCU-Train	Harp-Test	0.50 ± 0.29	0.50 ± 0.31	0.50 ± 0.29
Harp-Train + HCU-Train	Harp-Test	0.89 ± 0.01	0.89 ± 0.01	0.89 ± 0.02
Harp-Train + HCU-Train	HCU-Test	0.85 ± 0.04	0.69 ± 0.35	0.73 ± 0.33

Another experiment attempts to learn from both datasets at the same time (Table 4). The dataset now is the concatenation of HarP and HCUnicamp. The datasets were mixed together with a 70% training, 10% validation and 20% testing hold-out. The presence of patients and controls is balanced between the sets. Also included are results from testing in a different domain while training in other.

E2DHipseg was able to achieve good Dice in both the HarP and HCU when both are involved on training. However, while looking at only left or right results, poor Dice standard deviation is still present, meaning problems with resection are still happening. While examining predictions from training only in HCU and testing in HarP, in many cases the method predicted a resection was present in darker scans, when it wasn't, resulting in high false negatives.

### 5.3. Qualitative results

While visually inspecting HarP results, very low variance was found. We noted no presence of heavy outliers. Other methods present similar, stable results. (Fig. 4.)

However, in HCUnicamp, way more errors are visible in the worst segmentations in Fig. 5(b). Specially where the hippocampus is removed. Other methods have similar results, with false positives in voxels where the hippocampus would be in a healthy subject or Alzheimer's patient. As expected, the best segmentation, displayed in Fig. 5(a), was in a control, healthy subject.

## 6. Discussion

Regarding the Consensus approach from our method, most of the false positives some of the networks produce are eliminated by the averaging of activations followed by thresholding and post processing. This approach allows the methodology to focus on good segmentation on the hippocampus area, without worrying with small false positives in other areas of the brain. It was also observed that in some cases, one

of the networks fails and the consensus of the other two "saves" the result.

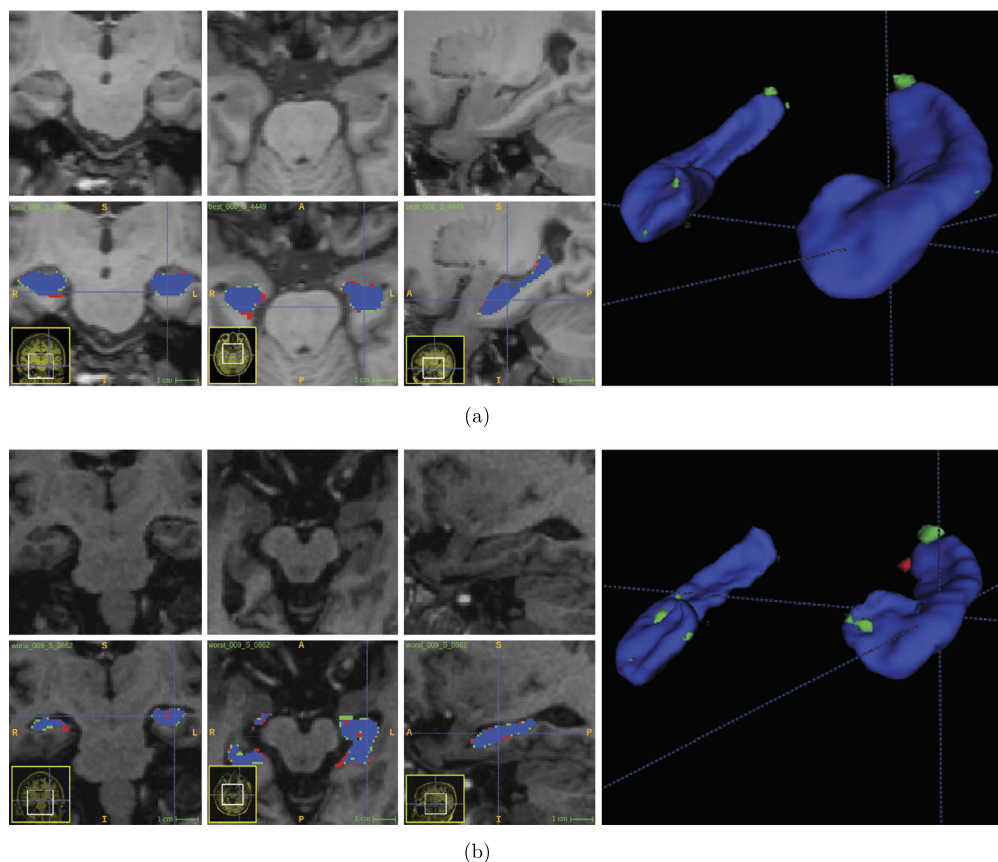
The fact that patches are randomly selected and augmented in runtime means they are mostly not repeated in different epochs. This is different to making a large dataset of pre-processed patches with augmentation. We believe this random variation during training is very important to ensure the network keeps seeing different data in different epochs, improving generalization. This idea is similar to the Dropout technique [35], only done in data instead of weights. Even with this patch randomness, re-runs of the same experiment resulted mostly in the same final results, within 0.01 mean Dice of each other.

As visible on the results of multiple methods, Dice when evaluating using the HCUnicamp dataset is not on the same level as what is seen on the public benchmark. Most methods have false positives on the removed hippocampus area, in a similar fashion to Fig. 5(b). The fact that QuickNat and Hippodeep have separate outputs for left and right hippocampus does not seem to be enough to solve this problem. We believe the high false positive rate is due to textures similar to the hippocampus, present in the hippocampus area, after its removal.

Final experiments attempt to adapt the methodology to Epilepsy volumes. Training in HCUnicamp improved results, but the high standard deviation and mistakes on hippocampus resections are still present. A similar story is seen while analyzing results from concatenating the HarP train and HCU-Train dataset in training. The method was able to achieve good overall Dice in both the HarP test set and HCU-Test, of 0.89 and 0.85, but while analyzing right and left hippocampus separately the high standard deviation due to missed resections was still present. The resulting mean Dice was low due to cases of false positives in resections on the left or right Dice resulting in 0 Dice, pulling the mean Dice down drastically. This was confirmed in the qualitative results and does not happen when training and testing in HCUnicamp controls or HarP, as showcased by the similar, low standard deviation between overall Dice and left/right Dice. This problem could possibly be solved with a preliminary hippocampus presence detection phase in future work, but this is not in the scope of this paper, since HCUnicamp was used here as a test set and this approach would be a bias to the test set.

## 7. Conclusion

This paper presents a hippocampus segmentation method including consensus of multiple U-Net based CNNs and traditional post-processing, successfully using a new optimizer and loss function from the literature. The presented method achieves state-of-the-art performance on the public HarP hippocampus segmentation benchmark. The hypothesis was raised that current automatic hippocampus segmentation methods, including our own, would not have the same performance on our in-house epilepsy dataset, with many cases of hippocampus removal. Quantitative and qualitative results show failure from those methods to take into account hippocampus removal, in unseen epilepsy



**Fig. 4.** Multiview and 3D render (approximate) of our (a) best and (b) worst cases while evaluating in the HarP test set. Prediction in green, target in red and overlap in purple.

data. This raises the concern that current automatic hippocampus segmentation methods are not ready to deal with hippocampus resection due to epilepsy treatment. We show that training in the epilepsy data does improve results, but there is still room for improvement. In future work, improvements can be made to our methodology to detect the removal of the hippocampus as a pre-processing step.

## Declarations

### Author contribution statement

Diedre Carmo: Conceived and designed the experiments; Performed the experiments; Analyzed and interpreted the data; Wrote the paper.

Bruna Silva: Contributed reagents, materials, analysis tools or data.

Clarissa Yasuda: Contributed reagents, materials, analysis tools or data; Wrote the paper.

Letícia Rittner, Roberto Lotufo: Conceived and designed the experiments; Wrote the paper.

### Funding statement

This work was supported by São Paulo Research Foundation (FAPESP) and CAPES under grant 2018/00186-0. CNPq also supported this work under grants 310828/2018-0 and 308311/2016-7.

### Data availability statement

Data associated with this study has been deposited at <http://www.hippocampal-protocol.net/SOPs/index.php>. The authors do not have permission to share data from HCUncamp.

### Declaration of interests statement

The authors declare no conflict of interest.

### Additional information

Supplementary content related to this article has been published online at <https://doi.org/10.1016/j.heliyon.2021.e06226>.

### Acknowledgements

Alzheimer's disease data collection and sharing for this project was provided 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](http://www.fnih.org)). The grantee organization is the

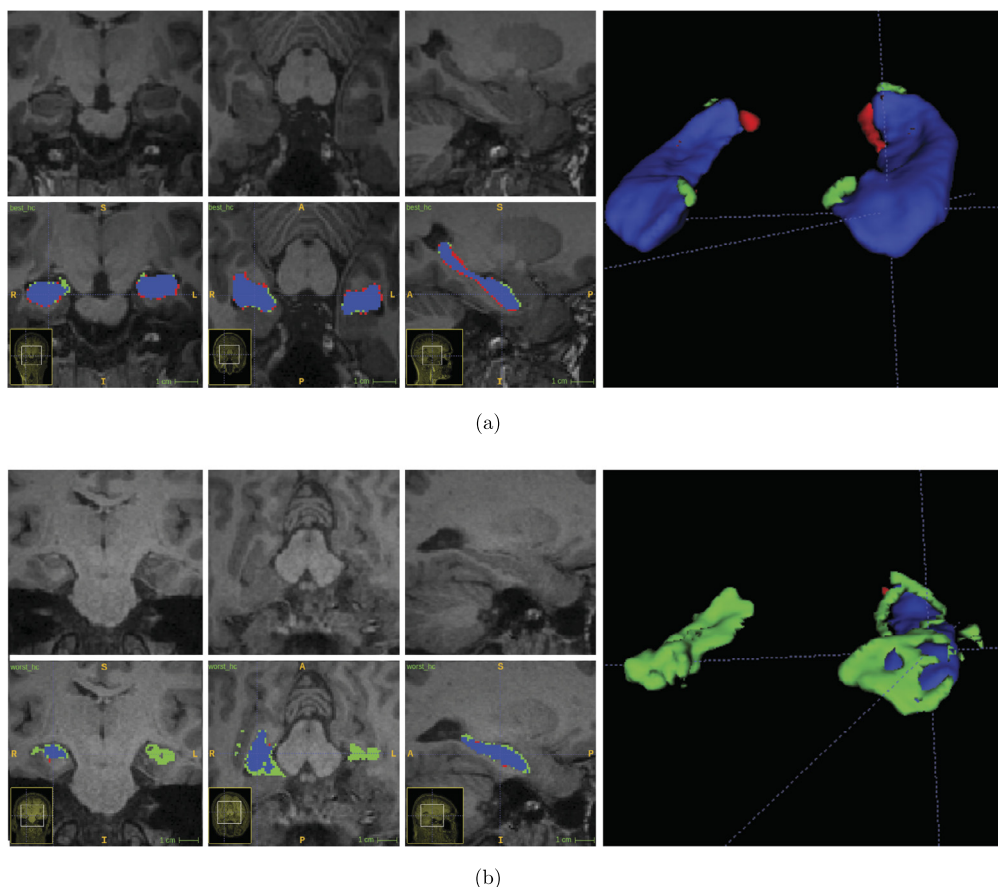


Fig. 5. Multiview and 3D render of our (a) best and (b) worst cases while testing in the HCUnicamp dataset. Prediction in green, target in red and overlap in purple.

Northern California Institute for Research and Education, and the study is coordinated by the Alzheimer's Therapeutic Research Institute at the University of Southern California. ADNI data are disseminated by the Laboratory for Neuro Imaging at the University of Southern California.

Additionally, we thank our partners at BRAINN for letting us use their epilepsy dataset on this research.

## References

- [1] P. Andersen, *The Hippocampus Book*, Oxford University Press, 2007.
- [2] R.C. Petersen, P. Aisen, L.A. Beckett, M. Donohue, A. Gamst, D.J. Harvey, C. Jack, W. Jagust, L. Shaw, A. Toga, et al., Alzheimer's disease neuroimaging initiative (ADNI): clinical characterization, *Neurology* 74 (3) (2010) 201–209.
- [3] E. Ghizoni, R.N. Matias, S. Lieber, B.M. de Campos, C.L. Yasuda, P.C. Pereira, A.C.S. Amato Filho, A.F. Joaquim, T.M. Lopes, H. Tedeschi, et al., Clinical and imaging evaluation of transcus selective amygdalohippocampotomy, *World Neurosurg.* 100 (2017) 665–674.
- [4] E. Ghizoni, J. Almeida, A.F. Joaquim, C.L. Yasuda, B.M. de Campos, H. Tedeschi, F. Cendes, Modified anterior temporal lobectomy: anatomical landmarks and operative technique, *J. Neurol. Surg. Part A: Cent. Eur. Neurosurg.* 76 (05) (2015) 407–414.
- [5] B. Fischl, *FreeSurfer*, *NeuroImage* 62 (2) (2012) 774–781.
- [6] C.S. McCarthy, A. Ramprasad, C. Thompson, J.-A. Botti, I.L. Coman, W.R. Kates, A comparison of FreeSurfer-generated data with and without manual intervention, *Front. Neurosci.* 9 (2015) 379.
- [7] A.G. Roy, S. Conjeti, N. Navab, C. Wachinger, A.D.N. Initiative, et al., QuickNAT: a fully convolutional network for quick and accurate segmentation of neuroanatomy, *NeuroImage* 186 (2019) 713–727.
- [8] B. Thyreau, K. Sato, H. Fukuda, Y. Taki, Segmentation of the hippocampus by transferring algorithmic knowledge for large cohort processing, *Med. Image Anal.* 43 (2018) 214–228.
- [9] F. Isensee, P. Kickingereder, W. Wick, M. Bendszus, K.H. Maier-Hein, Brain tumor segmentation and radiomics survival prediction: contribution to the BraTS 2017 challenge, in: *International MICCAI Brainlesion Workshop*, Springer, 2017, pp. 287–297.
- [10] H. Wang, J.W. Suh, S.R. Das, J.B. Pluta, C. Craige, P.A. Yushkevich, Multi-atlas segmentation with joint label fusion, *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (3) (2013) 611–623.
- [11] J.E. Iglesias, M.R. Sabuncu, Multi-atlas segmentation of biomedical images: a survey, *Med. Image Anal.* 24 (1) (2015) 205–219.
- [12] J. Pipitone, M.T.M. Park, J. Winterburn, T.A. Lett, J.P. Lerch, J.C. Pruessner, M. Lepage, A.N. Voineskos, M.M. Chakravarty, A.D.N. Initiative, et al., Multi-atlas segmentation of the whole hippocampus and subfields using multiple automatically generated templates, *NeuroImage* 101 (2014) 494–512.
- [13] A. Chincarini, F. Sensi, L. Rei, G. Gemme, S. Squarcia, R. Longo, F. Brun, S. Tangaro, R. Bellotti, N. Amoroso, et al., Integrating longitudinal information in hippocampal volume measurements for the early detection of Alzheimer's disease, *NeuroImage* 125 (2016) 834–847.
- [14] C. Platero, M.C. Tobar, Combining a patch-based approach with a non-rigid registration-based label fusion method for the hippocampal segmentation in Alzheimer's disease, *Neuroinformatics* 15 (2) (2017) 165–183.
- [15] Y. Chen, B. Shi, Z. Wang, P. Zhang, C.D. Smith, J. Liu, Hippocampus segmentation through multi-view ensemble ConvNets, in: *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, IEEE, 2017, pp. 192–196.
- [16] Z. Xie, D. Gillies, Near real-time hippocampus segmentation using patch-based canonical neural network, arXiv preprint, arXiv:1807.05482.
- [17] C. Wachinger, M. Reuter, T. Klein, DeepNAT: deep convolutional neural network for segmenting neuroanatomy, *NeuroImage* 170 (2018) 434–445.
- [18] D. Ataloglou, A. Dimou, D. Zarpalas, P. Daras, Fast and precise hippocampus segmentation through deep convolutional neural network ensembles and transfer learning, *Neuroinformatics* (2019) 1–20.
- [19] N.K. Dinsdale, M. Jenkinson, A.I. Namburete, Spatial warping network for 3D segmentation of the hippocampus in MR images, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2019, pp. 284–291.
- [20] C.H. Sudre, W. Li, T. Vercauteren, S. Ourselin, M.J. Cardoso, Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations, in: *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, Springer, 2017, pp. 240–248.
- [21] O. Ronneberger, P. Fischer, T. Brox, U-Net: convolutional networks for biomedical image segmentation, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2015, pp. 234–241.
- [22] A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet classification with deep convolutional neural networks, in: *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [23] N. Nogovitsyn, R. Souza, M. Muller, A. Srajer, S. Hassel, S.R. Arnott, A.D. Davis, G.B. Hall, J.K. Harris, M. Zamyadi, et al., Testing a deep convolutional neural net-



- work for automated hippocampus segmentation in a longitudinal sample of healthy participants, *NeuroImage* 197 (2019) 589–597.
- [24] M. Boccardi, M. Bocchetta, F.C. Morency, D.L. Collins, M. Nishikawa, R. Ganzola, M.J. Grothe, D. Wolf, A. Redolfi, M. Pievani, et al., Training labels for hippocampal segmentation based on the EADC-ADNI harmonized hippocampal protocol, *Alzheimer's Dement.* 11 (2) (2015) 175–183.
- [25] D. Carmo, B. Silva, C. Yasuda, L. Rittner, R. Lotufo, Extended 2D consensus hippocampus segmentation, arXiv preprint, arXiv:1902.04487.
- [26] O. Lucena, R. Souza, L. Rittner, R. Frayne, R. Lotufo, Silver standard masks for data augmentation applied to deep-learning-based skull-stripping, in: *Biomedical Imaging (ISBI 2018)*, 2018 IEEE 15th International Symposium on, IEEE, 2018, pp. 1114–1117.
- [27] M. Brett, K. Christoff, R. Cusack, J. Lancaster, et al., Using the Talairach atlas with the MNI template, *NeuroImage* 13 (6) (2001) 85–85.
- [28] M. Pereira, R. Lotufo, L. Rittner, An extended-2D CNN approach for diagnosis of Alzheimer's disease through structural MRI, in: *Proceedings of the 27th Annual Meeting of ISMRM 2019*, 2019, p. na.
- [29] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [30] S. Ioffe, C. Szegedy, Batch normalization: accelerating deep network training by reducing internal covariate shift, arXiv preprint, arXiv:1502.03167.
- [31] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, arXiv preprint, arXiv:1409.1556.
- [32] V. Iglovikov, A. Shvets, TeraNet: U-Net with VGG11 encoder pre-trained on ImageNet for image segmentation, arXiv preprint, arXiv:1801.05746.
- [33] H. Kervadec, J. Bouchtiba, C. Desrosiers, E. Granger, J. Dolz, I. Ben Ayed, Boundary loss for highly unbalanced segmentation, in: M.J. Cardoso, A. Feragen, B. Glocker, E. Konukoglu, I. Oguz, G. Unal, T. Vercauteren (Eds.), *Proceedings of the 2nd International Conference on Medical Imaging with Deep Learning*, in: *Proceedings of Machine Learning Research*, vol. 102, PMLR, London, United Kingdom, 2019, pp. 285–296. URL <http://proceedings.mlr.press/v102/kervadec19a.html>.
- [34] E.R. Dougherty, R.A. Lotufo, *Hands-on Morphological Image Processing*, vol. 59, SPIE Press, 2003.
- [35] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* 15 (1) (2014) 1929–1958.