Medical Image Analysis 48 (2018) 177–186

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Automated sub-cortical brain structure segmentation combining spatial and deep convolutional features

Kaisar Kushibar[1],*, Sergi Valverde[1], Sandra González-Villà, Jose Bernal, Mariano Cabezas, Arnau Oliver, Xavier Lladó

Institute of Computer Vision and Robotics, University of Girona, Ed. P-IV, Campus Montilivi, Girona, 17003, Spain

ARTICLE INFO

Article history:
Received 26 September 2017
Revised 1 March 2018
Accepted 9 June 2018
Available online 15 June 2018

Keywords:
Brain
MRI
Sub-cortical structures
Segmentation
Convolutional neural networks

ABSTRACT

Sub-cortical brain structure segmentation in Magnetic Resonance Images (MRI) has attracted the interest of the research community for a long time as morphological changes in these structures are related to different neurodegenerative disorders. However, manual segmentation of these structures can be tedious and prone to variability, highlighting the need for robust automated segmentation methods. In this paper, we present a novel convolutional neural network based approach for accurate segmentation of the sub-cortical brain structures that combines both convolutional and prior spatial features for improving the segmentation accuracy. In order to increase the accuracy of the automated segmentation, we propose to train the network using a restricted sample selection to force the network to learn the most difficult parts of the structures. We evaluate the accuracy of the proposed method on the public MICCAI 2012 challenge and IBSR 18 datasets, comparing it with different traditional and deep learning state-of-the-art methods. On the MICCAI 2012 dataset, our method shows an excellent performance comparable to the best participant strategy on the challenge, while performing significantly better than state-of-the-art techniques such as FreeSurfer and FIRST. On the IBSR 18 dataset, our method also exhibits a significant increase in the performance with respect to not only FreeSurfer and FIRST, but also comparable or better results than other recent deep learning approaches. Moreover, our experiments show that both the addition of the spatial priors and the restricted sampling strategy have a significant effect on the accuracy of the proposed method. In order to encourage the reproducibility and the use of the proposed method, a public version of our approach is available to download for the neuroimaging community.

1. Introduction

Brain structure segmentation in Magnetic Resonance Images (MRI) is one of the major interests in medical practice due to its various applications, including pre-operative evaluation and surgical planning, radiotherapy treatment planning, longitudinal monitoring for disease progression or remission (Kikinis et al., 1996; Phillips et al., 2015; Pitiot et al., 2004), etc. The sub-cortical structures (i.e. thalamus, caudate, putamen, pallidum, hippocampus, amygdala, and accumbens) have attracted the interest of the research community for a long time, since their morphological changes are frequently associated with psychiatric and neurodegenerative disorders and could be used as biomarkers of some diseases (Debernard et al., 2015; Mak et al., 2014). Therefore, segmentation of sub-cortical brain structures in MRI for quantitative analysis has a major clinical application. However, manual segmentation of MRI is extremely time consuming and hardly reproducible due to inter- and intra- variability among operators, highlighting the need for automated accurate segmentation methods.

Recently, González-Villà et al. (2016), reviewed different approaches for brain structure segmentation in MRI. One of the commonly used automatic brain structure segmentation tools in medical practice is FreeSurfer,[2] which uses non-linear registration and an atlas-based segmentation approach (Fischl et al., 2002). Another classical approach, also popular in the medical community,

---

* Corresponding author.
E-mail addresses: kaisar.kushibar@udg.edu (K. Kushibar), sergio.valverde@udg.edu (S. Valverde), sgonzalez@eia.udg.edu (S. González-Villà), jose.bernal@udg.edu (J. Bernal), mariano.cabezas@udg.edu (M. Cabezas), aoliver@eia.udg.edu (A. Oliver), xavier.llado@udg.edu (X. Lladó).

[1] These authors contributed equally to this work.

[2] https://surfer.nmr.mgh.harvard.edu/.

https://doi.org/10.1016/j.media.2018.06.006

is the method proposed by Patenaude et al. (2011) – FIRST, which is included into the publicly available software FSL.[3] This method uses the principles of Active Shape (Cootes et al., 1995) and Active Appearance Models (Cootes et al., 2001) that are put within a Bayesian framework, allowing to use the probabilistic relationship between shape and intensity to its full extent.

In recent years, deep learning methods, in particular, Convolutional Neural Networks (CNN), have demonstrated a state-of-the-art performance in many computer vision tasks such as visual object detection, classification and segmentation (Krizhevsky et al., 2012; He et al., 2016; Szegedy et al., 2015; Girshick et al., 2014). Unlike handcrafted features, CNN methods learn from observed data (LeCun et al., 1998) making relevant features to a specific task. Naturally, CNNs are also becoming a popular technique applied in medical image analysis. There have been many advances in the application of deep learning in medical imaging such as expert-level performance in skin cancer classification (Esteva et al., 2017), high rate detecting cancer metastases (Liu et al., 2017), Alzheimer's disease classification (Sarraf and Tofighi, 2016), and spotting early signs of autism (Hazlett et al., 2017).

Some CNN methods have also been proposed for brain structure segmentation. One of the common ways used in the literature is patch-based segmentation, where patches of a certain size are extracted around each voxel and classified using a CNN. Application of 2D, 3D, 2.5D patches (patches from the three orthogonal views of an MRI volume) and their combinations including multi-scale patches can be found in the literature for brain structure segmentation (Brébisson and Montana, 2015; Bao and Chung, 2016; Milletari, 2017; Mehta et al., 2017). Combining patches of different views and dimensions is done in a multi-path manner, where CNNs consist of different branches corresponding to each patch type, i.e. parallel interconnected processing modules analyze each of the inputs. In contrast to patch-based CNNs, fully convolutional neural networks (FCNN) produce segmentation for a neighborhood of an input patch (Long et al., 2015). Shakeri et al. (2016) adapted the work of Chen et al. (2016) for semantic segmentation of natural images using FCNN. Moreover, 3D FCNNs, which segment a 3D neighborhood of an input patch at once, have been investigated by Dolz et al. (2018) and Wachinger et al. (2018). Although FCNNs show improvement in segmentation speed due to parallel segmentation of several voxels, they suffer from a high number of parameters in the network in comparison with patch-based CNNs.

It is common to apply post-processing methods to refine the final segmentation output. Inference of CNN-priors and statistical models such as Markov Random Fields and Conditional Random Fields (Lafferty et al., 2001) were used in the experiments of Brébisson and Montana (2015), Shakeri et al. (2016), and Wachinger et al. (2018). A modified Random Walker based segmentation refinement has been also proposed by Bao and Chung (2016).

Apart from implicit information that is provided by the extracted patches from MRI volumes, explicit characteristics distinguishing spatial consistency have been studied. Brébisson and Montana (2015) included distances to centroids to their networks. Wachinger et al. (2018) used the Euclidean and spectral coordinates computed from eigenfunctions of a Laplace-Beltrami operator of a solid 3D brain mask, to provide a distinctive perception of spatial location for every voxel. These kinds of features provide additional spatial information, however, extracting these explicit features from an unannotated MRI volume requires some preliminary operations to be attended (e.g. repetitive training of the network to compute initial segmentation mask).

From the reviewed literature, we have observed that most of the current deep learning approaches for sub-cortical brain structure segmentation focus on segmenting only the large sub-cortical structures (thalamus, caudate, putamen, pallidum). However, other important small structures (i.e. hippocampus, amygdala, accumbens), which are used for examining neurological disorders such as schizophrenia (Altshuler et al., 1998; Lawrie et al., 2003), anxiety disorder (Milham et al., 2005), bipolar disorder (Altshuler et al., 1998), Alzheimer (Fox et al., 1996), etc., are not considered. These small structures have smaller volume – hence, lower number of samples – compared to the other larger structures, which hinders training deep learning strategies and makes the segmentation task more challenging. In this paper, we present our approach for segmenting the sub-cortical structures: a new 2.5D CNN architecture – i.e., the three orthogonal views of a 3D volume – that incorporates probabilistic atlases as spatial features. Although probabilistic atlases have been used before in deep learning methods (Ghafoorian et al., 2017), they have never been applied for segmenting the sub-cortical brain structures. Within our research, unlike most of the existing deep learning approaches, we address segmenting all the sub-cortical structures, including the smallest ones. To the best of our knowledge, this is the first deep learning method incorporating atlas probabilities into a CNN for sub-cortical brain structure segmentation. Moreover, we propose a particular sample selection technique, which allows the neural network to learn to segment the most difficult areas of the structures in the images, and also show its importance in achieving higher accuracy. We test the proposed strategy in two well-known datasets: MICCAI 2012[4] (Landman and Warfield, 2012) and IBSR 18[5]; and compare our results with the classical and recent CNN strategies for brain structure segmentation. Additionally, we make our method publicly available for the community, accessible online at https://github.com/NIC-VICOROB/sub-cortical_segmentation.

2. Method

2.1. Input features

In our method, we employ 2.5D patches to incorporate information from three orthogonal views of a 3D volume. In our case, each patch has a size of 32×32 pixels. Although 3D patches may provide more information of surroundings for the voxel that is being classified, they are computationally and memory expensive. Thus, by using 2.5D patches, we approximate the information that is provided by a 3D patch in computational time and memory efficient manner.

Along with the appearance based features provided by the T1-w MRI, we employ spatial features extracted from a structural probabilistic atlas. In our experiments, we used the well-known Harvard–Oxford (Caviness et al., 1996) atlas template in MNI152 space distributed with the FSL package,[6] which has been built using 47 young adult healthy brains. In our method, first, T1-w image of the MNI152 template is affine registered to T1-w image of the considered datasets using a block matching approach (Ourselin et al., 2000). Then, non-linear registration of the atlas template to subject volume is applied using fast free-form deformation method (Modat et al., 2010). The deformation field obtained after the registration is used to move the probabilistic atlas into the subject space. Registration processes have been carried out using the well known and publicly available tool NiftyReg.[7] Afterwards, vectors of size 15, corresponding to seven anatomical structures with left and right parts separately and background, were

---

[3] https://fsl.fmrib.ox.ac.uk/fsl/fslwiki.

[4] https://masi.vuse.vanderbilt.edu/workshop2012.

[5] https://www.nitrc.org/projects/ibsr.

[6] https://fsl.fmrib.ox.ac.uk/fsl/fslwiki.

[7] http://cmictig.cs.ucl.ac.uk/wiki/index.php/NiftyReg.