



Brain tumor segmentation with Deep Neural Networks



Mohammad Havaei^{a,*}, Axel Davy^b, David Warde-Farley^c, Antoine Biard^{c,d}, Aaron Courville^c, Yoshua Bengio^c, Chris Pal^{c,e}, Pierre-Marc Jodoin^a, Hugo Larochelle^a

^a Université de Sherbrooke, Sherbrooke, Qc, Canada

^b École Normale supérieure, Paris, France

^c Université de Montréal, Montréal, Canada

^d École polytechnique, Palaiseau, France

^e École Polytechnique de Montréal, Canada

ARTICLE INFO

Article history:

Received 27 April 2015

Revised 2 March 2016

Accepted 11 May 2016

Available online 19 May 2016

Keywords:

Brain tumor segmentation

Deep neural networks

Convolutional neural networks

Cascaded convolutional neural networks

ABSTRACT

In this paper, we present a fully automatic brain tumor segmentation method based on Deep Neural Networks (DNNs). The proposed networks are tailored to glioblastomas (both low and high grade) pictured in MR images. By their very nature, these tumors can appear anywhere in the brain and have almost any kind of shape, size, and contrast. These reasons motivate our exploration of a machine learning solution that exploits a flexible, high capacity DNN while being extremely efficient. Here, we give a description of different model choices that we've found to be necessary for obtaining competitive performance. We explore in particular different architectures based on Convolutional Neural Networks (CNN), i.e. DNNs specifically adapted to image data.

We present a novel CNN architecture which differs from those traditionally used in computer vision. Our CNN exploits both local features as well as more global contextual features simultaneously. Also, different from most traditional uses of CNNs, our networks use a final layer that is a convolutional implementation of a fully connected layer which allows a 40 fold speed up. We also describe a 2-phase training procedure that allows us to tackle difficulties related to the imbalance of tumor labels. Finally, we explore a cascade architecture in which the output of a basic CNN is treated as an additional source of information for a subsequent CNN. Results reported on the 2013 BRATS test data-set reveal that our architecture improves over the currently published state-of-the-art while being over 30 times faster.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

In the United States alone, it is estimated that 23,000 new cases of brain cancer will be diagnosed in 2015.¹ While gliomas are the most common brain tumors, they can be less aggressive (i.e. low grade) in a patient with a life expectancy of several years, or more aggressive (i.e. high grade) in a patient with a life expectancy of at most 2 years.

Although surgery is the most common treatment for brain tumors, radiation and chemotherapy may be used to slow the growth of tumors that cannot be physically removed. Magnetic resonance imaging (MRI) provides detailed images of the brain, and is one of the most common tests used to diagnose brain tumors. All the more, brain tumor segmentation from MR images can have great

impact for improved diagnostics, growth rate prediction and treatment planning.

While some tumors such as meningiomas can be easily segmented, others like gliomas and glioblastomas are much more difficult to localize. These tumors (together with their surrounding edema) are often diffused, poorly contrasted, and extend tentacle-like structures that make them difficult to segment. Another fundamental difficulty with segmenting brain tumors is that they can appear anywhere in the brain, in almost any shape and size. Furthermore, unlike images derived from X-ray computed tomography (CT) scans, the scale of voxel values in MR images is not standardized. Depending on the type of MR machine used (1.5, 3 or 7 tesla) and the acquisition protocol (field of view value, voxel resolution, gradient strength, b0 value, etc.), the same tumorous cells may end up having drastically different gray-scale values when pictured in different hospitals.

Healthy brains are typically made of 3 types of tissues: the white matter, the gray matter, and the cerebrospinal fluid. The goal of brain tumor segmentation is to detect the location and

* Corresponding author.

E-mail address: seyed.mohammad.havaei@usherbrooke.ca (M. Havaei).

¹ <http://www.ucancer.org>

extension of the tumor regions, namely active tumorous tissue (vascularized or not), necrotic tissue, and edema (swelling near the tumor). This is done by identifying abnormal areas when compared to normal tissue. Since glioblastomas are infiltrative tumors, their borders are often fuzzy and hard to distinguish from healthy tissues. As a solution, more than one MRI modality is often employed, e.g. T1 (spin-lattice relaxation), T1-contrasted (T1C), T2 (spin-spin relaxation), proton density (PD) contrast imaging, diffusion MRI (dMRI), and fluid attenuation inversion recovery (FLAIR) pulse sequences. The contrast between these modalities gives almost a unique signature to each tissue type.

Most automatic brain tumor segmentation methods use hand-designed features (Farahani et al., 2014; Menze et al., 2014). These methods implement a classical machine learning pipeline according to which features are first extracted and then given to a classifier whose training procedure does not affect the nature of those features. An alternative approach for designing task-adapted feature representations is to *learn* a hierarchy of increasingly complex features directly from in-domain data. Deep neural networks have been shown to excel at learning such feature hierarchies (Bengio et al., 2013). In this work, we apply this approach to learn feature hierarchies adapted specifically to the task of brain tumor segmentation that combine information across MRI modalities.

Specifically, we investigate several choices for training CNNs, which are DNNs adapted to image data. We report their advantages, disadvantages and performance using well established metrics. Although CNNs first appeared over two decades ago (LeCun et al., 1998), they have recently become a mainstay of the computer vision community due to their record-shattering performance in the ImageNet Large-Scale Visual Recognition Challenge (Krizhevsky et al., 2012). While CNNs have also been successfully applied to segmentation problems (Alvarez et al., 2012; Long et al., 2015; Hariharan et al., 2014; Ciresan et al., 2012), most of the previous work has focused on non-medical tasks and many involve architectures that are not well suited to medical imagery or brain tumor segmentation in particular. Our preliminary work on using convolutional neural networks for brain tumor segmentation together with two other methods using CNNs was presented in BRATS'14 workshop. However, those results were incomplete and required more investigation (More on this in Section 2).

In this paper, we propose a number of specific CNN architectures for tackling brain tumor segmentation. Our architectures exploit the most recent advances in CNN design and training techniques, such as Max-out (Goodfellow et al., 2013b) hidden units and Dropout (Srivastava et al., 2014) regularization. We also investigate several architectures which take into account both the local shape of tumors as well as their context.

One problem with many machine learning methods is that they perform pixel classification without taking into account the local dependencies of labels (i.e. segmentation labels are conditionally independent given the input image). To account for this, one can employ structured output methods such as conditional random fields (CRFs), for which inference can be computationally expensive. Alternatively, one can model label dependencies by considering the pixel-wise probability estimates of an initial CNN as additional input to certain layers of a second DNN, forming a cascaded architecture. Since convolutions are efficient operations, this approach can be significantly faster than implementing a CRF.

We focus our experimental analysis on the fully-annotated MICCAI brain tumor segmentation (BRATS) challenge 2013 data-set (Farahani et al., 2014) using the well defined training and testing splits, thereby allowing us to compare directly and quantitatively to a wide variety of other methods.

Our contributions in this work are four fold:

1. We propose a fully automatic method with results currently ranked second on the BRATS 2013 scoreboard;
2. To segment a brain, our method takes between 25 s and 3 min, which is one order of magnitude faster than most state-of-the-art methods.
3. Our CNN implements a novel two-pathway architecture that learns about the local details of the brain as well as the larger context. We also propose a two-phase training procedure which we have found is critical to deal with imbalanced label distributions. Details of these contributions are described in Sections 3.1.1 and 3.2.
4. We employ a novel cascaded architecture as an efficient and conceptually clean alternative to popular structured output methods. Details on those models are presented in Section 3.1.2.

2. Related work

As noted by Menze et al. (2014), the number of publications devoted to automated brain tumor segmentation has grown exponentially in the last several decades. This observation not only underlines the need for automatic brain tumor segmentation tools, but also shows that research in that area is still a work in progress.

Brain tumor segmentation methods (especially those devoted to MRI) can be roughly divided in two categories: those based on generative models and those based on discriminative models (Menze et al., 2014; Bauer et al., 2013; Angelini et al., 2007).

Generative models rely heavily on domain-specific prior knowledge about the appearance of both healthy and tumorous tissues. Tissue appearance is challenging to characterize, and existing generative models usually identify a tumor as being a shape or a signal which deviates from a normal (or average) brain (Clark et al., 1998). Typically, these methods rely on anatomical models obtained after aligning the 3D MR image on an atlas or a template computed from several healthy brains (Doyle et al., 2013). A typical generative model of MR brain images can be found in Prastawa et al. (2004). Given the ICBM brain atlas, the method aligns the brain to the atlas and computes posterior probabilities of healthy tissues (white matter, gray matter and cerebrospinal fluid). Tumorous regions are then found by localizing voxels whose posterior probability is below a certain threshold. A post-processing step is then applied to ensure good spatial regularity. Prastawa et al. (2003), also register brain images onto an atlas in order to get a probability map for abnormalities. An active contour is then initialized on this map and iterated until the change in posterior probability is below a certain threshold. Many other active-contour methods along the same lines have been proposed (Khotanlou et al., 2009; Cobzas et al., 2007; Popuri et al., 2012), all of which depend on left-right brain symmetry features and/or alignment-based features. Note that since aligning a brain with a large tumor onto a template can be challenging, some methods perform registration and tumor segmentation at the same time (Kwon et al., 2014; Parisot et al., 2012).

Other approaches for brain tumor segmentation employ discriminative models. Unlike generative modeling approaches, these approaches exploit little prior knowledge on the brain's anatomy and instead rely mostly on the extraction of [a large number of] low level image features, directly modeling the relationship between these features and the label of a given voxel. These features may be raw input pixels values (Havaei et al., 2014; Hamamci et al., 2012), local histograms (Kleesiek et al., 2014; R.Meier et al., 2014) texture features such as Gabor filterbanks (Subbanna et al., 2013; 2014), or alignment-based features such as inter-image gradient, region shape difference, and symmetry analysis (N.Tustison and Avants, 2013). Classical discriminative learning techniques such as SVMs (Bauer et al., 2011; Schmidt et al., 2005; Lee et al., 2005) and decision forests (Zikic et al., 2012) have also been used. Results

Download English Version:

<https://daneshyari.com/en/article/6878002>

Download Persian Version:

<https://daneshyari.com/article/6878002>

[Daneshyari.com](https://daneshyari.com)