University of Dundee

Fully Connected Crf With Data-driven Prior for Multi-class Brain Tumor Segmentation

Shen, Haocheng; Zhang, Jianguo

Published in:
2017 IEEE International Conference on Image Processing (ICIP)

DOI:
[10.1109/ICIP.2017.8296577](https://doi.org/10.1109/ICIP.2017.8296577)

Publication date:
2018

Document Version
Peer reviewed version

[Link to publication in Discovery Research Portal](#)

Citation for published version (APA):
Shen, H., & Zhang, J. (2018). Fully Connected Crf With Data-driven Prior for Multi-class Brain Tumor Segmentation. In *2017 IEEE International Conference on Image Processing (ICIP): 17-20 Sept. 2017* (pp. 1727-1731). (IEEE Conference Proceedings). IEEE. <https://doi.org/10.1109/ICIP.2017.8296577>

General rights

Copyright and moral rights for the publications made accessible in Discovery Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from Discovery Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain.
- You may freely distribute the URL identifying the publication in the public portal.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

FULLY CONNECTED CRF WITH DATA-DRIVEN PRIOR FOR MULTI-CLASS BRAIN TUMOR SEGMENTATION

Haocheng Shen and Jianguo Zhang

CVIP, Computing, School of Science and Engineering, University of Dundee, UK

ABSTRACT

Grid conditional random fields (CRFs) are widely applied in both natural and medical image segmentation tasks. However, they only consider the label coherence in neighborhood pixels or regions, which limits their ability to model long-range connections within the image and generally results in excessive smoothing of tumor boundaries. In this paper, we present a novel method for brain tumor segmentation in MR images based on fully-connected CRF (FC-CRF) model that establishes pairwise potentials on all pairs of pixels in the images. We employ a hierarchical approach to differentiate different structures of tumor and further formulate a FC-CRF model with learned data-driven prior knowledge of tumor core. The methods were evaluated on the testing and leaderboard set of Brain Tumor Image Segmentation Benchmark (BRATS) 2013 challenge. The precision of segmented tumor boundaries is improved significantly and the results are competitive compared to the start-of-the-arts.

Index Terms— CRF, prior, brain tumor segmentation

1. INTRODUCTION

Gliomas are the most frequent primary brain tumors in adults [1] and can be classified as high-grade(HG) or low-grade (LG) based on the aggressive form of the disease. Fig. 1 shows a typical gliomas tumor in different modalities and expert's annotation. Normally there are four structures in the tumor: edema(green), necrosis(red), non-enhancing tumor(blue) and enhancing tumor(yellow). The last three also make up a super-structure called *tumor core*. Automatic segmentation of gliomas brain tumor is an active topic for decade with challenges on the diversity and variation of tumor size, shape, and location and appearance.

A common approach is to pose this problem as classifying pixels into different tissues, followed by a CRF model which incorporates smoothness terms that maximize label agreement between pixels in the neighborhood [1, 2, 3]. However, the resulting adjacency CRF is limited in its ability to model long-range connections within the image and generally results in excessive smoothing of tumor boundaries. Clinically, the precisely segmented boundaries of tumors are crucial for making treatment plans as well as guiding surgery. Ignoring

Fig. 1: (a) Tumor structures in different modalities: edema(green), necrosis(red), non-enhancing tumor(blue) and enhancing tumor(yellow). Best viewed in colour. (b) The probability curve of being tumor core with respect to the distance to the centroid of the complete tumor region.

tumor boundaries may cause irreversible impact, for example, the loss of brain functionality such as speaking or reading. In the fully connected CRF, each node is assumed to be connected with every other. Thus it is able to consider not only short-range but also long-range interactions between pixels. Efficient inference for fully connected (FC)-CRF was first proposed by [4] for multi-class natural image segmentation and it has been recently applied into medical image domain [5, 6, 7]. Especially, [7] applied FC-CRF to brain tumor segmentation but it only refined the boundaries between complete tumor and background. Also a fair comparison between fully-connected and grid CRF is missing.

In this paper, we propose an automatic method for brain tumor segmentation based on FC-CRF. The main contributions of our paper include: 1) unlike [7], we apply FC-CRF for multi-class brain tumour segmentation (complete tumor and tumor core) through a hierarchical approach; 2) we conducted a fair comparison between FC and grid CRF and demonstrated that FC-CRF *significantly* improved the tumour boundary segmentation accuracy with very small p-value; 3) our approach is ranked top among automated methods on BRATS 2013 testing set; 4) we introduce a prior-driven FC-CRF and evaluated the effect of data-driven prior on different sets and demonstrate the prior is helpful on larger datasets.

2. METHODS

We use a *hierarchical* approach to segment the tumor tissues. In the first layer, we differentiate the non-tumor and tumor tissues using FC-CRF model. On the resulting complete tumor region, we apply the second layer to distinguish edema and tumor core using another FC-CRF. The final layer further partitions the segmented tumor core into necrosis, enhancing and non-enhancing using classifier only as visually the necrosis could be some scatter points in the tumor core and does not hold the assumption of local label consistency.

2.1. Fully-Connected CRF

Consider $\mathbf{y} = \{y_i\}$ as a labeling defined over all voxels of the 3D image in the domain of labels $L = \{0, 1\}$. The corresponding Gibbs energy function over this label configuration is composed of unary energy ψ_u and pairwise energy ψ_p :

$$E(\mathbf{y}) = \sum_i \psi_u(y_i, \mathbf{x}_i) + \sum_{i < j} \psi_p(y_i, y_j, \mathbf{f}_i, \mathbf{f}_j) \quad (1)$$

where $\mathbf{x}_i$ and $(\mathbf{f}_i, \mathbf{f}_j)$ are the unary and pairwise features, respectively. The unary potentials $\psi_u(y_i, \mathbf{x}_i)$ is computed independently for each voxel by a classifier that produces a probability distribution over the labels given the feature. The unary feature used in our implementation are described in Section 3.1. Pairwise potentials defines a similarity function of voxel features and labels to model their interactions. We use a form of pairwise potential similar to that in [4]:

$$\psi_p(y_i, y_j, \mathbf{f}_i, \mathbf{f}_j) = \mu(y_i, y_j) \sum_{m=1}^M w^{(m)} k^{(m)}(\mathbf{f}_i, \mathbf{f}_j) \quad (2)$$

where each $k^{(m)}(\cdot, \cdot)$ is a Gaussian kernel of an feature pair $(\mathbf{f}_i, \mathbf{f}_j)$, $w^{(m)}$ is a linear combination weight, and $\mu(y_i, y_j)$ is a label compatibility function. Our pairwise kernels have the following form:

$$k(\mathbf{f}_i, \mathbf{f}_j) = w^{(1)} \exp\left(-\frac{|\mathbf{p}_i - \mathbf{p}_j|^2}{2\theta_p^2} - \frac{|\mathbf{f}_i - \mathbf{f}_j|^2}{2\theta_f^2}\right) + w^{(2)} \exp\left(-\frac{|\mathbf{p}_i - \mathbf{p}_j|^2}{2\theta_\alpha^2}\right) \quad (3)$$

where $\mathbf{p}_i$ and $\mathbf{p}_j$ are the 3D coordinate vector of pixel i and j . We use voxel intensities from different modalities as the feature vector $\mathbf{f}_i$ and $\mathbf{f}_j$. The degree of nearness and feature similarity are controlled by parameter θ_p and θ_f . The second Gaussian kernel term is used to smooth the segmentation results and is controlled by parameter θ_α . The compatibility function μ is given by the Potts model, $\mu(y_i, y_j) = [y_i \neq y_j]$. It penalizes nearby similar pixels that are assigned different labels.

Note that the pairwise potentials are a linear combination of Gaussian kernels over a Euclidean feature space. We implemented a 3D efficient inference approach on a permutohedral lattice using bilateral filtering method [4, 8], which could be completed in a linear time with respect to the number of nodes rather than quadratic time.

2.2. Data-driven Prior for Tumor Core FC-CRF

One drawback of fully connected CRF is that it is sensitive to the initialization based on the unary potential. In our case, if the output from classifier is vague, it is challenge for CRF to correct large regions of misclassified pixels (as shown in Fig. 4(c)). In this section, we propose to utilize data-driven prior with unary potential to improve the initialization of fully connected CRF when distinguishing tumor core against edema in the second layer of hierarchy.

We observed that although tumor can occur almost everywhere in the brain, one topology of tumor structures is explicit: the tumor core is usually surrounded by edema region. Motivated by this observation, we build a probability model to interpret the prior for the progressive level of the tumor core in the complete tumor region. Such a probability model is learned directly on the training set. Specifically, we calculate the probability P_p representing the core progressive level at voxel i with respect to its distance d to the centroid of the complete tumor region using the experts' annotations. For each complete tumor region T_j , we first estimate $P_j(y_i|d) = N_{y_i}/N_d$ where N_{y_i} is the total number of voxels labeled as $y_i \in \{1, -1\}$ (1 for tumor core and -1 for edema) within distance d : $N_{y_i} = \#\{i|y_i, \gamma_j \|i - c_j\|_2 \leq d, i \in T_j\}$ where c_j is the centroid of tumor T_j . N_d is the total number of voxels within distance d : $N_d = \#\{i|\gamma_j \|i - c_j\|_2 \leq d, i \in T_j\}$. The final P_p is calculated based all of the n tumors, defined as $P_p(y_i|d) = \frac{\sum_{j=1}^n P_j(y_i|d)}{\sum_{j=1}^n \sum_{y_i} P_j(y_i|d)} = 1/n \sum_{j=1}^n P_j(y_i|d)$. Considering the size of complete tumor region varies across different patients, the distance d within each tumor is normalized into $[0, 1]$ using $\gamma_j = \max_{i \in T_j} \|i - c_j\|_2$, i.e., invariant to scaling. As P_p is learned from training data, we call it data-driven prior. To the best of our knowledge, such a prior has not been explored before in MR brain tumor segmentation.

Fig.1(b) plots the learned probabilities of a voxel being tumor core at different distances. The curves show a strong progressive prior for tumor core in both HG and LG images: the closer to the centroid, the higher probability the voxel being tumor core. The data-driven prior is further fitted with a sigmoid function. When the voxel is further to the centroid, its probability prior is more likely affected by the tumor boundaries often of irregular shape. Therefore, we only calculate and use the prior probability for voxels within the ball with a radius of $(\mu_j + 2 * \sigma_j)$ where μ_j and σ_j are the mean and standard deviation of all distances within complete tumor T_j . To incorporate this prior, the unary potential in the second

layer CRF is reformulated as:

$$\psi_u(y_i, \mathbf{x}_i) = -\log(w * P_c(y_i|\mathbf{x}_i) + (1-w) * P_p(y_i|d)) \quad (4)$$

where $P_c(y_i|\mathbf{x}_i)$ is the probability output of classifier and w is the trade-off between classifier output and prior knowledge.

3. EVALUATION

We evaluated the proposed method on BRATS 2013 clinical dataset. The dataset is comprised of 3 sub-datasets: 30 patients subjects (20 HGs and 10 LGs) with pixel-level annotations for the training set; 10 (all HGs) for the testing set and 25 patient subjects (21 HGs and 4 LGs) for the leaderboard set. For each subject there exists 4 modalities, namely T1, T1-contrast (T1c), T2 and Flair which are all skull-stripped and co-registered.

Quantitative evaluation on the test and leaderboard set is through the on-line VSD evaluation system¹ for three sub-tasks: 1) the *complete* tumor region (including all four tumor structures); 2) the *core* tumor region (including all tumor structure except "edema"); 3) the *enhancing* tumor region (including only the "enhancing tumor" structure). For each tumor region, *Dice*, *Sensitivity* (*Se*) and *Positive Predictive Value* (*PPV*) are computed. Due to space limitation, we only report *Dice* here as it is the most common metric for segmentation task. The full metric tables are provided in the supplemental file.

3.1. Implementation

For pre-processing, we first clip the intensity values to the quantile range [0.01, 0.99] and then rescale intensities to [0, 1]. Finally, we normalize intensities of each brain MR image with zero mean and unit standard deviation.

We choose the following unary features in our implementation: 1) intensity; 2) the hemispheric intensity difference between two symmetric pixels in the axial view. The symmetric axis in the every axial slice is computed using the approach presented in [9]; 3) first order statistics in a $5 \times 5 \times 5$ volume patch; and 4) maximum response filter (MR8) [10]. We extend the MR8 filter bank into 3D version by calculating the original MR8 filter responses in sagittal, axial and coronal views separately, and then concatenate the responses together.

The above features are calculated for each voxel and each modality. The features of all modalities are then concatenated into the final 128-dimensional feature vector for each voxel. We choose random forest as classifier.

3.2. Cross Validation

We compare FC and grid CRF on the training set using cross validation. We train different models for HG and LG: 5-fold

Table 1: Results of FC CRF and grid CRF in HG and LG images.

Methods	Dice		
	Complete	Core	Enhancing
Grid (HG)	0.85	0.71	0.72
FC (HG)	0.86	0.72	0.74
FC + Prior (HG)	0.86	0.74	0.73
Grid (LG)	0.86	0.63	N/A
FC (LG)	0.88	0.67	N/A
FC + Prior (LG)	0.88	0.70	N/A

Fig. 2: Segmentation accuracy of complete tumor boundaries w.r.t percentage of misclassified pixels within trimaps of different widths. (a) HG cases; (b) LG cases.

cross validation for 20 HG images and leave-one-out for 10 LG images. For grid CRF, we use the graph cuts toolbox from [11]. All the parameters $\{\theta_p, \theta_f, \theta_\alpha, \mathbf{w}\}$ are learned by grid search. The performance of FC and grid CRF are shown in Table 1. It is clear that FC-CRF outperforms grid CRF in both HG and LG images in all sub-tasks. To evaluate segmentation accuracy around boundaries, we adopt the "trimap" metric for the complete tumors as the same as in [4]. Specifically, we count the relative number of misclassified pixels within a narrow band ("trimap") surrounding actual tumor boundaries, obtained from the experts' ground truth. As shown in Fig. 2, the FC-CRF outperforms grid CRF across all trimap widths. We further conducted Wilcoxon signed-rank test on the results and reported the p-values for HG and LG cases, i.e., both values are $p = 7.8 \times 10^{-3}$, indicating that the improvements are statistically *significant* in both cases. Two sample results are shown in Fig. 3. It is obvious that FC-CRF generates more accurate boundaries while grid CRF over-segments.

Using the model with data-driven prior (Eq.4) the performance of tumor core segmentation can further increase. Fig. 4(c)&(d) visualize some segmented regions of tumor core by different models. Compared to the results without using prior, it is obvious that prior-driven FC-CRF is capable of not only correcting misclassified voxels in core regions but also removing some false positives.

¹<https://www.virtualskeleton.ch/>

Fig. 3: Visualization of result of grid and FC-CRF for complete tumor region: (a) Flair modality; (b) ground truth; (c) FC-CRF; (d) Grid CRF.

Fig. 4: Visualization of results of FC-CRF and FC-CRF with prior for tumor core (red) against edema (green): (a) T1c modality; (b) ground truth; (c) FC-CRF; (d) FC-CRF with prior. Best viewed in color.

3.3. Comparison with Best Performers

We then compare the proposed methods with currently published state-of-the-arts on BRATS-2013 testing set. As it only contains HG images, we only use the 20 HG training images for training. The results are listed in Table 2. The method of FC-CRF are in the top rank. Specifically, Tustison, Meier and Reza are the best performers of BRATS 2013 challenge [1]. Our method outperforms them all. Particularly, the winner of BRATS 2013 challenge [2] is less efficient than ours as it needs an auxiliary health brain dataset to do registration to calculate the asymmetric feature while we only use the data provided by the challenge. Our model is also slightly better than [12] but worse than [13]. They both use the popular Deep CNN model which is difficult to train without a modern GPU. Note that although [13] performs best on this dataset, they evaluated different experimental settings directly on the

Table 2: Comparison with the state-of-the-arts on testing set.

Method	Dice		
	Complete	Core	Enhancing
Pereira [13]	0.88	0.83	0.77
FC-CRF(ours)	0.88	0.78	0.74
Tustison [2]	0.87	0.78	0.74
Havaei [12]	0.88	0.78	0.73
FC-CRF+Prior(ours)	0.88	0.75	0.74
Meier [1]	0.82	0.73	0.69
Reza [1]	0.83	0.72	0.72

Table 3: Comparison with the state-of-the-arts on leaderboard.

Method	Dice		
	Complete	Core	Enhancing
Pereira [13]	0.84	0.72	0.62
Havaei [12]	0.84	0.71	0.57
FC-CRF+prior(ours)	0.81	0.65	0.55
FC-CRF(ours)	0.81	0.63	0.53
Tustison [2]	0.79	0.65	0.53
Reza [1]	0.72	0.60	0.53

testing set, which might lead to the model overfitted on the testing set. It is shown that using prior knowledge decreases the performance. As the testing images are relatively easier compared to the training images [1], one possible reason is the learned trade-off w in Eq.4 is overweighted where using classifier output can already generate good enough results.

The leaderboard results are listed in Table 3. This set is larger and relatively more challenge than testing set [1]. We observe improvements of using data-driven prior of tumor core. Our results are competitive with the state-of-the-arts except the deep CNN model [12, 13], and are better than all other shallow methods.

4. CONCLUSION

We proposed an automatic brain tumor segmentation method based on fully connected CRF. We further learned a prior of tumor cores and introduced a prior-driven FC-CRF for segmenting tumor cores. Though on the small set, prior does not seem to improve performance, it shows positive and promising effect on the larger datasets from which the results are more statistically conclusive and generalizable. Our results are competitive with the start-of-the-arts in BRATS 2013 testing and leaderboard set, and the segmented boundaries are significantly improved. In the future, we will explore CNN features within our framework.

Acknowledgments

This work was supported partially by the National Natural Science Foundation of China (No. 61628212).

5. REFERENCES

- [1] Bjoern H Menze et al., “The multimodal brain tumor image segmentation benchmark (BRATS),” *IEEE Transactions on Medical Imaging*, vol. 34, no. 10, pp. 1993–2024, 2015.
- [2] Nicholas J Tustison, KL Shrinidhi, Max Wintermark, et al., “Optimal symmetric multimodal templates and concatenated random forests for supervised brain tumor segmentation (simplified) with ANTsR,” *Neuroinformatics*, vol. 13, no. 2, pp. 209–225, 2015.
- [3] Nagesh Subbanna, Doina Precup, and Tal Arbel, “Iterative multilevel MRF leveraging context and voxel information for brain tumour segmentation in MRI,” in *CVPR*, 2014, pp. 400–405.
- [4] Philipp Krähenbühl and Vladlen Koltun, “Efficient inference in fully connected CRFs with gaussian edge potentials,” in *NIPS*, 2011, pp. 109–117.
- [5] Samuel Kadoury, Nadine Abi-Jaoudeh, and Pablo A Valdes, “Higher-order CRF tumor segmentation with discriminant manifold potentials,” in *MICCAI*, 2013, pp. 719–726.
- [6] José Ignacio Orlando and Matthew Blaschko, “Learning fully-connected CRFs for blood vessel segmentation in retinal images,” in *MICCAI*, 2014, pp. 634–641.
- [7] Konstantinos Kamnitsas, Christian Ledig, Virginia FJ Newcombe, Joanna P Simpson, Andrew D Kane, David K Menon, Daniel Rueckert, and Ben Glocker, “Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation,” *Medical Image Analysis*, vol. 36, pp. 61–78, 2017.
- [8] Andrew Adams, Jongmin Baek, and Myers Abraham Davis, “Fast high-dimensional filtering using the permutohedral lattice,” in *Computer Graphics Forum*, 2010, vol. 29, pp. 753–762.
- [9] Gareth Loy and Jan-Olof Eklundh, “Detecting symmetry and symmetric constellations of features,” in *ECCV*, 2006, pp. 508–521.
- [10] Manik Varma and Andrew Zisserman, “A statistical approach to texture classification from single images,” *International Journal of Computer Vision*, vol. 62, no. 1-2, pp. 61–81, 2005.
- [11] Brian Fulkerson, Andrea Vedaldi, and Stefano Soatto, “Class segmentation and object localization with superpixel neighborhoods,” in *ICCV*, 2009, pp. 670–677.
- [12] Mohammad Havaei, Axel Davy, David Warde-Farley, Antoine Biard, Aaron Courville, Yoshua Bengio, Chris Pal, Pierre-Marc Jodoin, and Hugo Larochelle, “Brain tumor segmentation with deep neural networks,” *Medical image analysis*, vol. 35, pp. 18–31, 2017.
- [13] Sérgio Pereira, Adriano Pinto, Victor Alves, and Carlos A Silva, “Brain tumor segmentation using convolutional neural networks in MRI images,” *IEEE transactions on Medical Imaging*, vol. 35, no. 5, pp. 1240–1251, 2016.