Simultaneous brain tumor and organs-at-risk segmentation for radiotherapy

Mikael Agn, DTU Compute

Introduction

When planning for radiotherapy of brain tumors several structures need to be segmented from MR images, such as tumors and so-called organs-at-risk (structures to be spared from radiation). Automatic segmentation is challenging since tumor location, shape, appearance and effect on healthy tissue vary greatly across patients. Moreover, MR intensities can vary significantly across scanners.

We propose a fully automated generative method for simultaneous segmentation of brain tumors and organs-at-risk (see fig. 2), We combine an existing whole-brain segmentation technique [1] with a spatial tumor prior, using convolutional restricted Boltzmann machines (cRBMs). The cRBMs are trained on expert tumor segmentations, without the use of any intensity information.

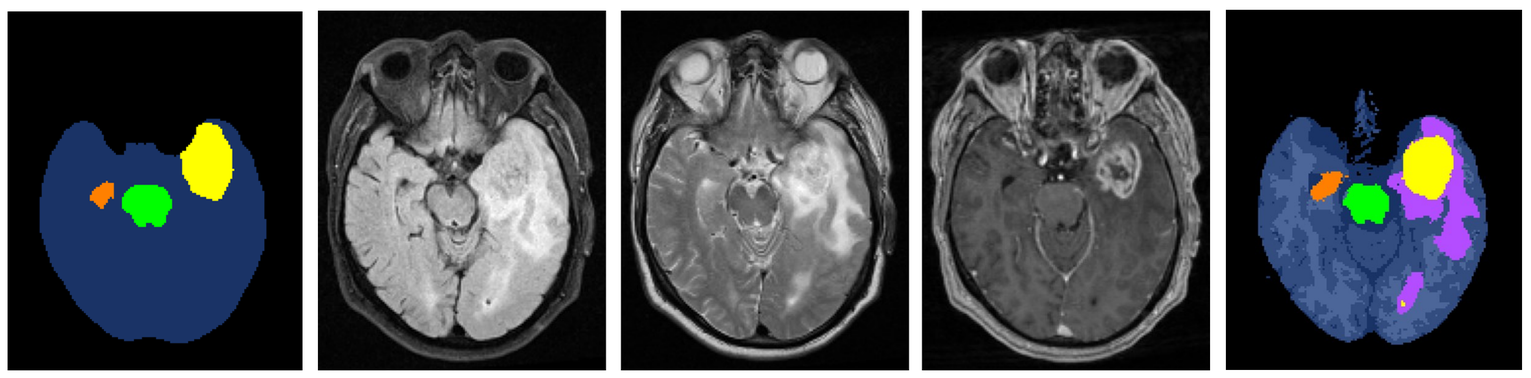

Fig 1: MR slices of a subject: Expert segmentation, Flair, T2, T1 with contrast and automatic segmentation. Automatic segmentation: Healthy labels $\mathbf{l}$ in dark blue, orange and green. Complete tumor label $\mathbf{z}$ in lilac and yellow. Core label $\mathbf{y}$ in yellow, existing inside $\mathbf{z}$.

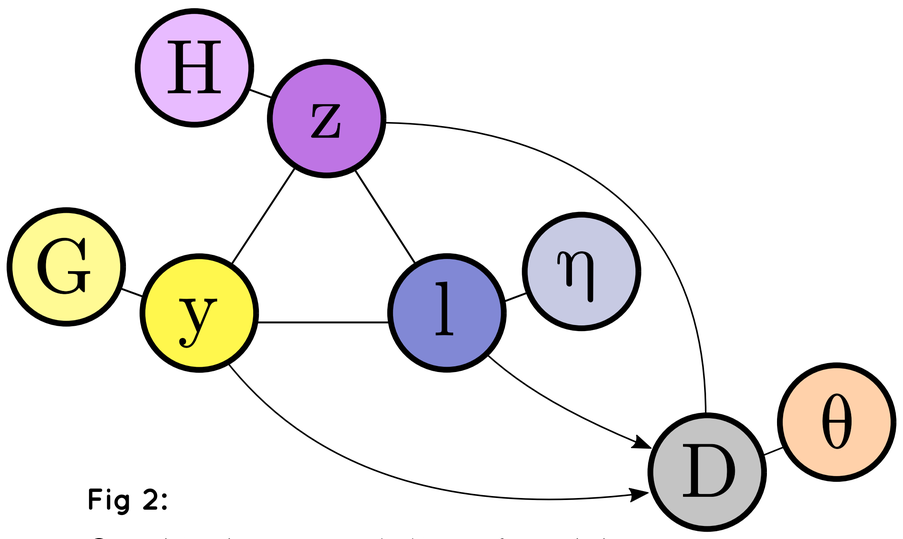

Fig 2:
Graphical representation of model.

Model

RBM tumor shape prior

- A convolutional restricted Boltzmann machine is a graphical model over visible and hidden units [2]. The visible and hidden layers are connected through convolutional filters $\mathbf{w}_k$.

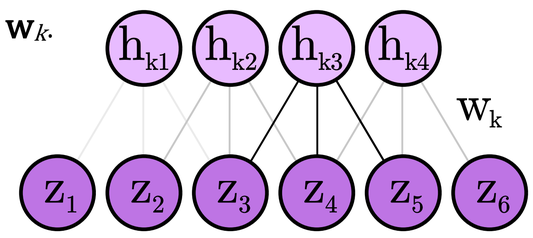

- We separately train one cRBM for the complete tumor label $\mathbf{z}$, with hidden units in $\mathbf{H}$, and one RBM for tumor core label $\mathbf{y}$, with hidden units $\mathbf{G}$.
- After training, the two cRBMs are combined to form a combined prior on tumor shape:

$$p(\mathbf{z}, \mathbf{y}) \propto \sum_{\mathbf{H},\mathbf{G}} \exp(-E_z(\mathbf{z}, \mathbf{H}) - E_y(\mathbf{y}, \mathbf{G}) - f(\mathbf{y}, \mathbf{z})),$$

where $f(\mathbf{y}, \mathbf{z})$ insures that core only exist within complete tumor and E denotes the cRBM energy term.

Likelihood function

- The likelihood function links labels to MR intensities in data $\mathbf{D}$. Each label is connected to a Gaussian mixture model with model parameters in $\boldsymbol{\theta}$. The likelihood is given by

$$p(\mathbf{D} \mid \mathbf{l}, \mathbf{z}, \mathbf{y}) = \int_{\boldsymbol{\theta}} p(\mathbf{D} \mid \mathbf{l}, \mathbf{z}, \mathbf{y}, \boldsymbol{\theta})\ p(\boldsymbol{\theta})\ d\boldsymbol{\theta}.$$

- Bias fields corrupting the MR images are modeled as linear combinations of spatially smooth basis functions, with parameters in $\boldsymbol{\theta}$.

Healthy tissue prior

For the prior on the healthy tissue labels in $\mathbf{l}$ we use a deformable probabilistic atlas computed from healthy subjects, consisting of mesh nodes $\boldsymbol{\eta}$ [3]. The healthy prior is given by

$$p(\mathbf{l}) = \int_{\boldsymbol{\eta}} p(\mathbf{l} \mid \boldsymbol{\eta})\, p(\boldsymbol{\eta})\, d\boldsymbol{\eta} \quad \text{and}$$

$$p(\mathbf{l} \mid \boldsymbol{\eta}) = \prod_i p(l_i \mid \boldsymbol{\eta}).$$

Inference

- **Step 1:** use $p(\mathbf{l}, \mathbf{z}, \mathbf{y} \mid \mathbf{D}) \simeq p(\mathbf{l}, \mathbf{z}, \mathbf{y} \mid \mathbf{D}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}})$, where $\langle \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\eta}} \rangle$ maximize $p(\boldsymbol{\theta}, \boldsymbol{\eta} \mid \mathbf{D})$.
  - Alternate between optimizing $\boldsymbol{\theta}$ by generalized Expectation-Maximization while keeping fixed $\boldsymbol{\eta}$, and $\boldsymbol{\eta}$ by conjugate gradient while keeping $\boldsymbol{\theta}$ fixed.
  - RBM energy temporarily replaced in this step with a simpler factorizable energy term.
- **Step 2:** Monte Carlo sampling from $p(\mathbf{l}, \mathbf{z}, \mathbf{y} \mid \mathbf{D}, \hat{\boldsymbol{\eta}})$, with $\boldsymbol{\eta}$ kept fixed.
  - Initialize with $\hat{\boldsymbol{\theta}}$.
  - Block-Gibbs sampling from $p(\mathbf{l}, \mathbf{z}, \mathbf{y}, \mathbf{H}, \mathbf{G}, \boldsymbol{\theta} \mid \mathbf{D}, \hat{\boldsymbol{\eta}})$, only retaining samples of $\mathbf{l}, \mathbf{z}$ and $\mathbf{y}$.

Experiments

- Each RBM was trained with 40 filters (size 7x7x7) on expert segmentations of 30 subjects for 9600 gradient steps of 0.1 with CD-1 and enhanced gradient [4].
- The method has been tested on 20 patients that have undergone radiotherapy treatment at Rigshospitalet.

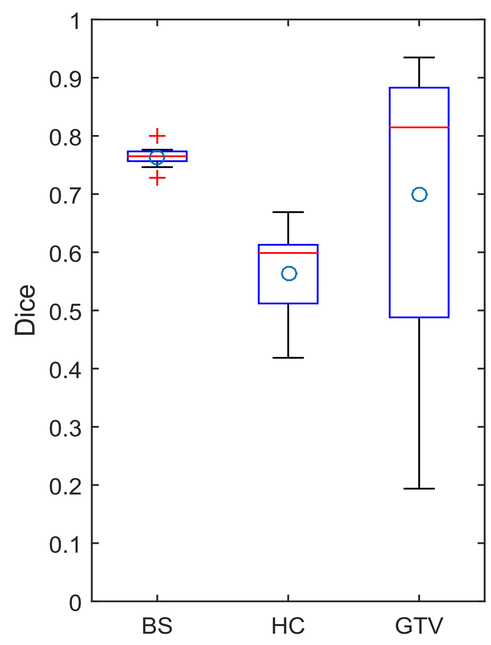

Fig 3: Box plot of Dice scores for brainstem (BS), hippocampus (HC) and gross tumor volume (GTV). Circles show mean values and central lines show median.

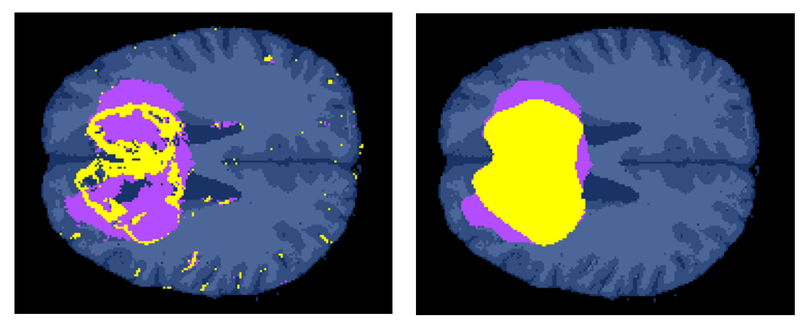

Fig 4: Automatic segmentation after step 1 on the left and final segmentation on the right.

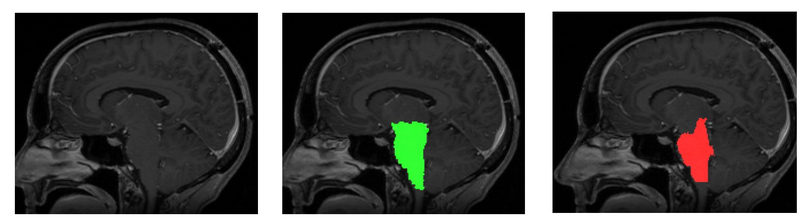

Fig 5: Brainstem segmentation. Manual in green, automatic in red. Note the difference in protocol.

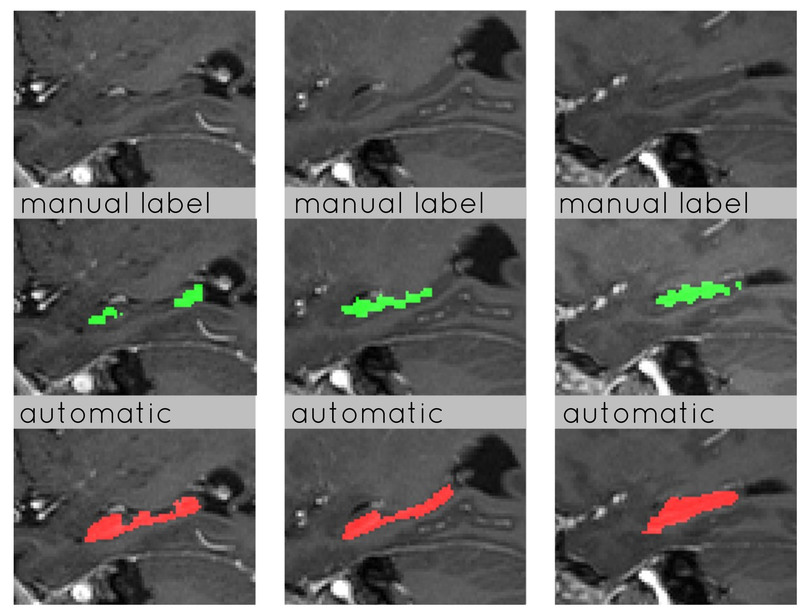

Fig 6: Hippocampus segmentations. Note that manual labels are less precise.

References

[1] Puonti, O., et al.: Fast, Sequence Adaptive Parcellation of Brain MR Using Parametric Models. In: Proc. MICCAI 2013. (2013)

[2]Lee, H., et al.: Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In: Proceedings of the 26th Annual International Conference on Machine Learning, ACM (2009)

[3] Van Leemput, K.: Encoding Probabilistic Brain Atlases Using Bayesian Inference. IEEE Trans Med Imag 28(6), (2009)

[4] Melchior, J., et al.: How to Center Binary Restricted Boltzmann Machines. arXiv preprint arXiv:1311.1354 (2013)