Lesion Detection Using Conditional Restricted Boltzmann Machines

Oula Puonti, DTU Compute

Introduction

- White matter lesions tend to occur in specific spatial context, for example around the ventricles (fig. 1)
- Previous lesion detection methods have used hand-crafted features [1] or random features [2]
- The **aim** of this project is to use Conditional Restricted Boltzmann Machines (cRBMs) to learn suitable features automatically, and use them for predicting where the lesions occur given a healthy brain anatomy
- The features are learned from manual segmentations done by radiologists

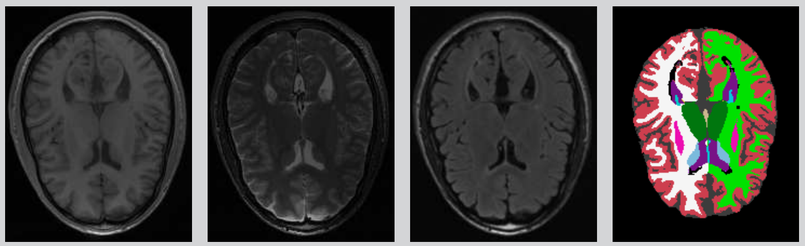

Figure 1: Three examples of MR scans with lesions and a manual segmentation.

Model Definition

- Model the conditional distribution of lesions given healthy tissue configuration using cRBMs [3]
- Given the model configuration (fig. 2) the distribution takes the form: $p(\mathbf{z}|\mathbf{l}) = \sum_{\mathbf{h}} \exp(-E(\mathbf{l}, \mathbf{z}, \mathbf{h}))/Z$
- The energy is defined as: $E(\mathbf{l}, \mathbf{z}, \mathbf{h}) = -\mathbf{l}^T\mathbf{W}_1\mathbf{h} - \mathbf{z}^T\mathbf{W}_2\mathbf{h} - \mathbf{l}^T\mathbf{W}_3\mathbf{z} - \mathbf{z}^T\mathbf{d} - \mathbf{h}^T\mathbf{b}$

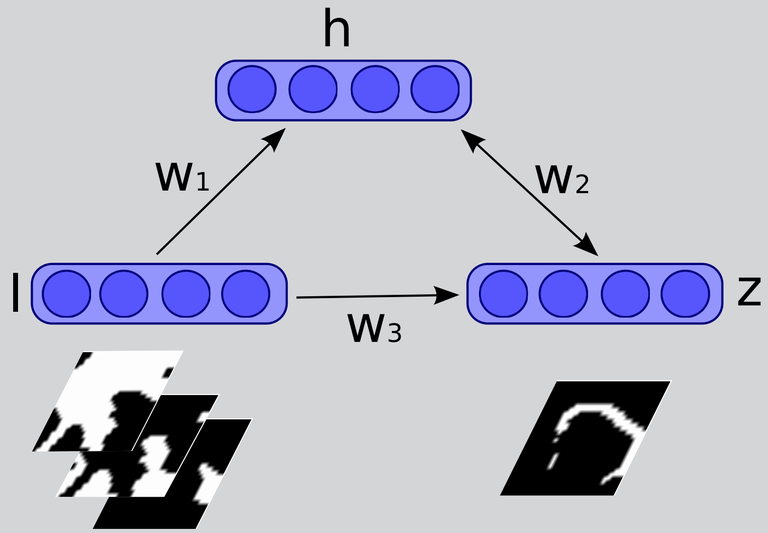

Figure 2: Model configuration.

Experiments

- The model was trained by minimizing the so-called generalized perceptron loss [3]
- Number of hidden units was 100 and the model was trained for 150 epochs with a learning rate of 1e-3
- One-of-K-coding was used for the manual segmentations
- The training data consists of 800 anatomy and simulated lesion patch pairs of size 30×30 (fig. 3)
- The learned features mostly rely on white matter and ventricle locations for predictions (fig. 4)
- **Exact inference** with the conditional is intractable: we use mean field approximation and sampling (fig. 5)

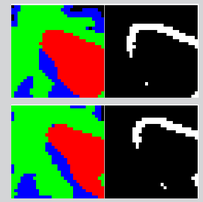

Figure 3: Simulated training data

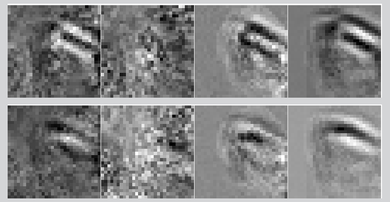

Figure 4: Learned features

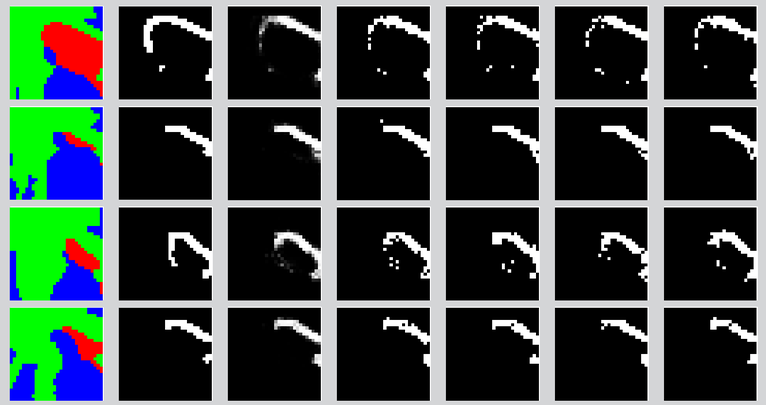

Figure 5: Test examples. From the left: true anatomy, true lesion, mean field approximation and 3 samples

Conclusions

- The model seems to learn reasonable features and is able to predict the lesion locations quite well
- Next step: how to go from individual patches up to full image size?

[1] Shiee et al. *A topology-preserving approach to the segmentation of brain images with multiple sclerosis lesions*, 2010

[2] Geremia et al. *Spatial Decision Forests for MS Lesion Segmentation in Multi-Channel MR Images*, 2010

[3] Mnih et al. *Conditional Restricted Boltzmann Machines for Structured Output Prediction*, 2012