Learning features for tissue classification with the classification restricted Boltzmann machine

Gijs van Tulder and Marleen de Bruijne

Motivation

Feature learners that learn efficient representations of their training data often ignore tissue class labels and may not be optimal for class discrimination. We investigate how the representation can be optimised for classification.

Restricted Boltzmann machines

An RBM is a probabilistic graphical model of its input and a hidden representation of that input. The classification RBM (Larochelle et al., 2012) extends the standard RBM with labels:

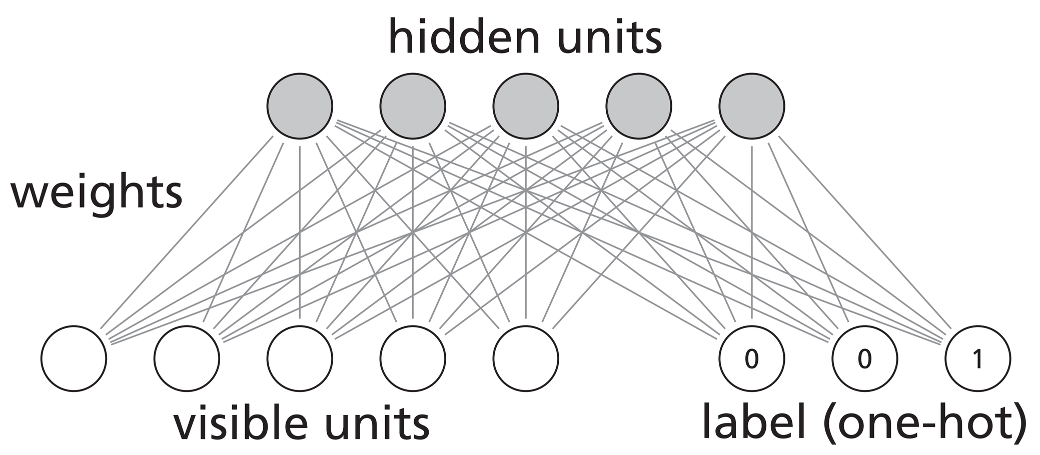

Learning objectives

The standard RBM learning objective optimises the probability of the inputs. In the classification RBM it optimises the visible units and the label. This is the generative learning objective

$$\log P(v_t, y_t).$$

The generative objective forces the RBM to model the largest variations in the data, which might be intra-class variation and not useful to a classifier.

The discriminative learning objective (Larochelle et al., 2012) maximises the posterior probability of the labels:

$$\log P(y_t \mid v_t).$$

This objective favours inter-class variability and is more likely to learn features that are helpful for classification.

We used a weighted combination of the two learning objectives:

$$\beta \log P(v_t, y_t) + (1 - \beta) \log P(y_t \mid v_t).$$

Two lung CT datasets

We use RBMs to learn features for two datasets. The first consists of 40 scans from the Danish Lung Cancer Screening Trial, with annotated airways. We detect airway patches.

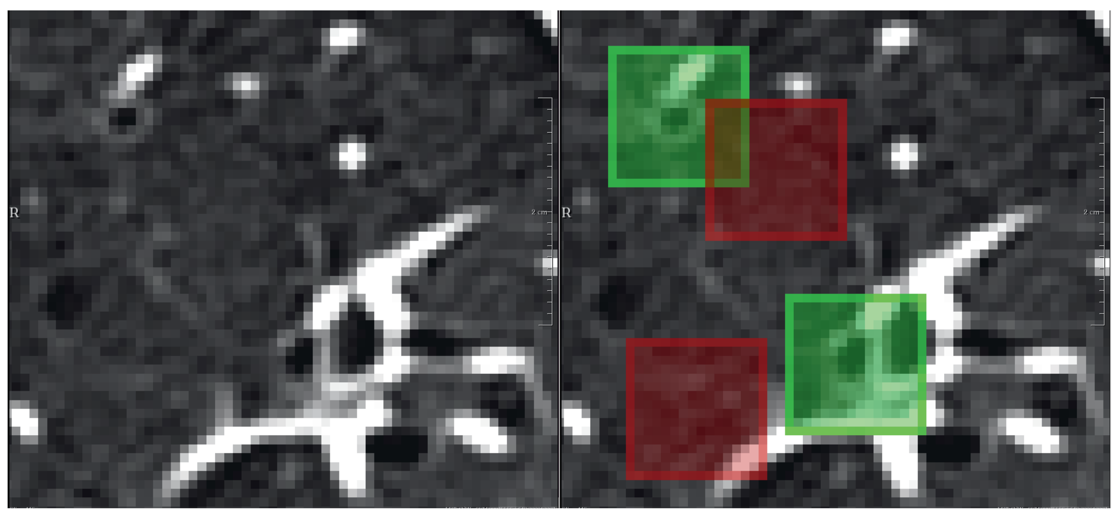

We also experiment with 73 scans from a dataset on interstitial lung diseases (Depeursinge et al., 2012). We use a convolutional RBM to classify subpatches into five lung tissue types.

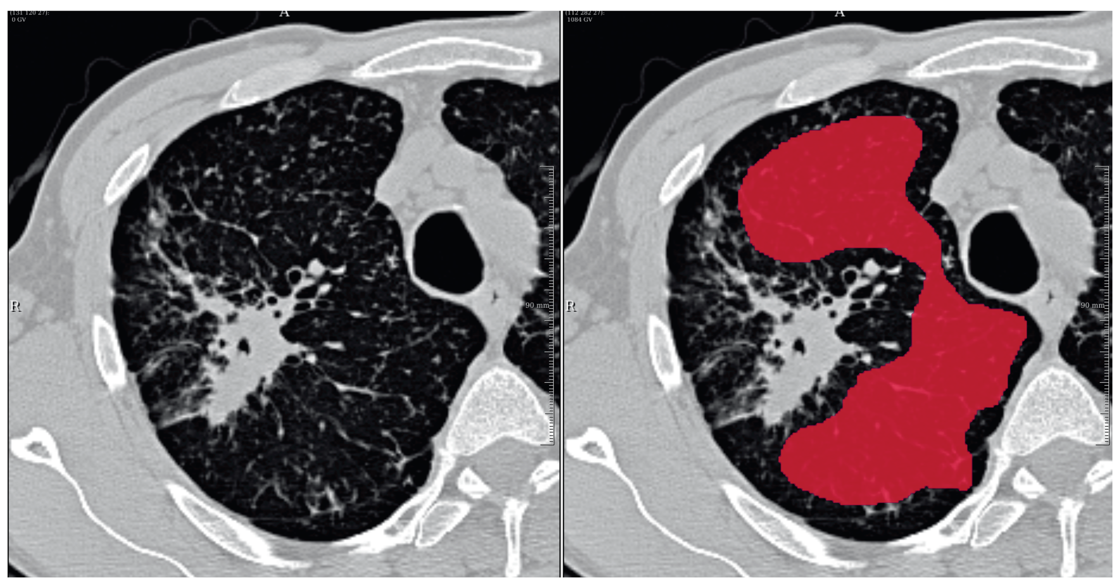

Classification accuracy

We found that a combination of generative and discriminative learning objectives often gave better results than using only one of the two objectives.

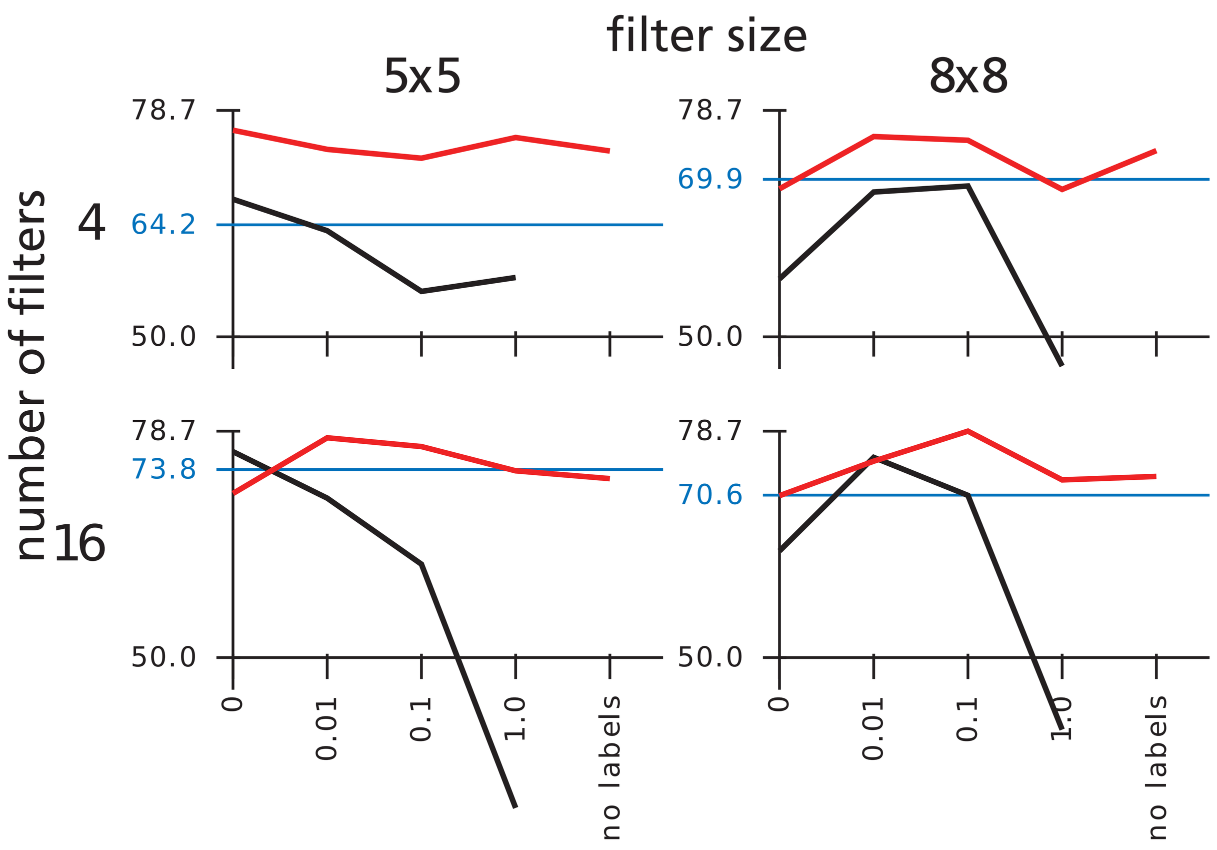

Test accuracy of the RBM (black), and of linear SVMs (red) trained on RBM-learned features. Blue lines show SVMs with random filters. The x-axis shows the amount of generative learning (β=0 is only discriminative, β=1 is only generative).

Number of filters

Using discriminative learning seems to be most important when the number of filters is small, when it is necessary to focus. The effect is less strong if the number of filters is large.

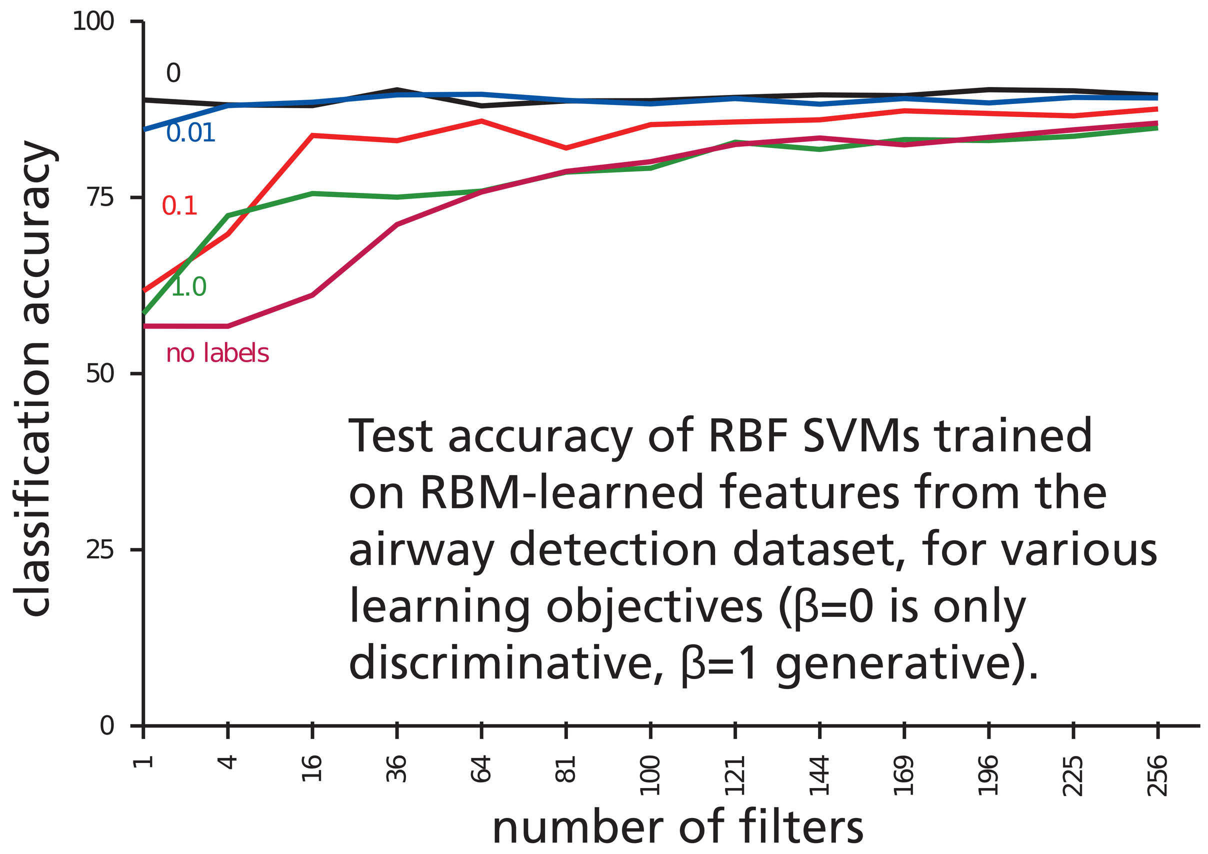

Test accuracy of RBF SVMs trained on RBM-learned features from the airway detection dataset, for various learning objectives (β=0 is only discriminative, β=1 generative).

Filter structure

The filters learned with a mixture of generative and discriminative learning showed more recognisable structure than the features learned with the purely discriminative objective.

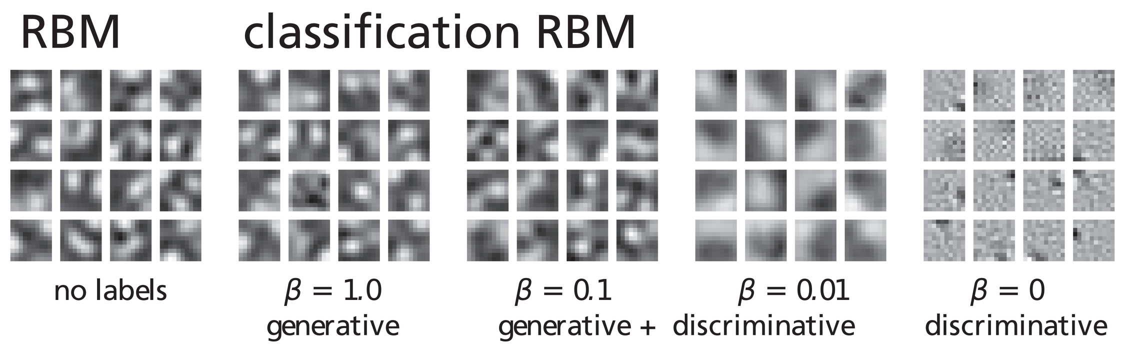

Example filters learned from the lung tissue data.

Conclusion

By training a classification RBM with a hybrid discriminative and generative learning objective, it can learn features that give a better classification result than features learned with only the generative learning objective.

References

Adrien Depeursinge et al. (2012). Building a reference multimedia database for interstitial lung diseases.
Hugo Larochelle et al. (2012). Learning Algorithms for the Classification Restricted Boltzmann Machine.

This research is financed by the Netherlands Organization for Scientific Research (NWO).

Biomedical Imaging Group Rotterdam, the Netherlands
www.bigr.nl · g.vantulder@erasmusmc.nl

MICCAI Workshop on Medical Computer Vision
September 2014, Boston, Massachusetts, USA

Erasmus MC
University Medical Center Rotterdam