Discriminative feature learning with restricted Boltzmann machines for lung tissue classification

Gijs van Tulder and Marleen de Bruijne

Summary

Standard RBMs are unsupervised, generative models. Their generative learning objective may produce the most faithful representation of the input, but it is not necessarily the most useful representation to use in a classifier.

We experimented with restricted Boltzmann machines (RBMs) for representation learning in medical image analysis. By training a classification RBM with a discriminative learning objective, we hoped that it would learn features that are better for classification. We found that a combination of generative and discriminative learning objectives gives features that give the best classification results.

Restricted Boltzmann machines

An RBM is a probabilistic graphical model of its input and a hidden representation of that input. The standard RBM is an unsupervised, generative model and does not include label information. It can be trained with a generative learning objective.

The classification RBM (Larochelle et al., 2012) extends the standard RBM with label information. This means it can be used for classification and can be trained with a discriminative learning objective.

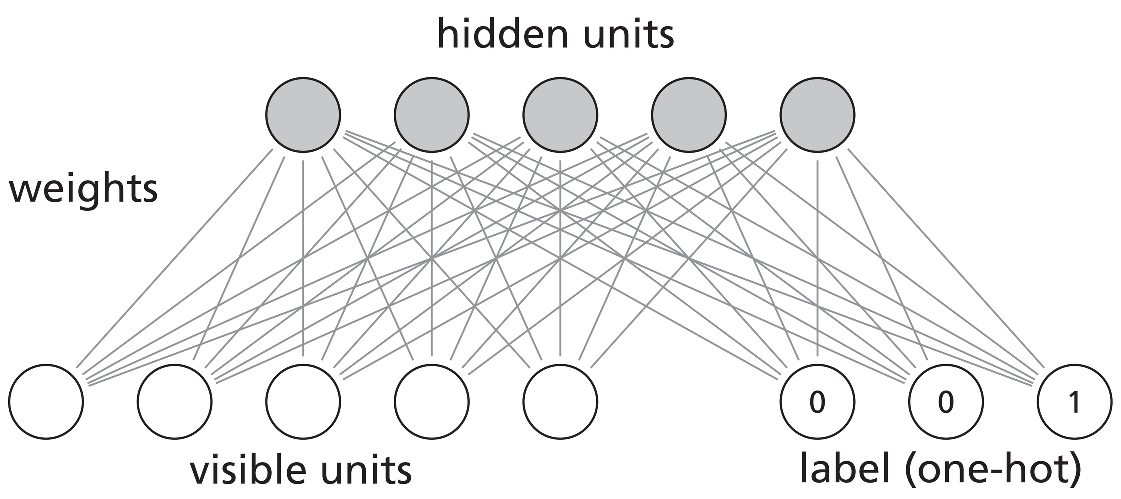

The classification RBM.

RBMs are usually trained with stochastic gradient descent and contrastive divergence.

Generative learning

The standard RBM learning objective tries to minimise the energy of the training samples of the model. In a classification RBM, this optimises the joint probability of the input and the labels, i.e., the generative objective

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

With this objective, the RBM models the strongest variations in the data, because these contribute most of the representation error.

Discriminative learning

If the representation is used as input for a classifier, the generative learning objective might waste effort on modelling the intra-class variation in the data, which is not useful to a classifier. Instead, we could use a learning objective that favours inter-class variability and features that are useful for classification.

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

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

In our experiments, we use a combination of the two learning objectives:

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

Lung tissue classification

We used RBM-based feature learning to build classifiers for lung tissue classification. In a dataset with subpatches taken from 73 lung scans, we classify each patch as one of four types of lung issue: healthy tissue or one of three disease patterns.

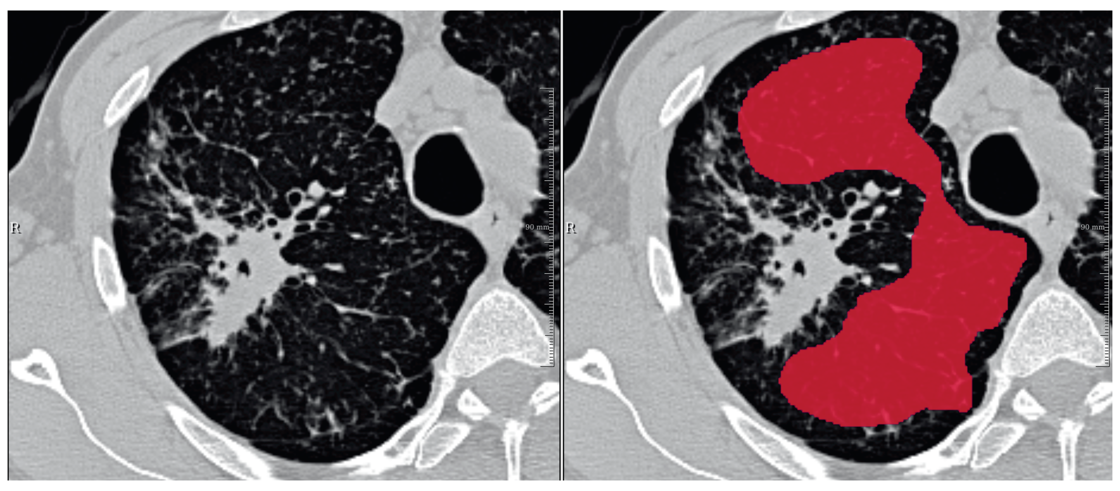

One slice from the lung tissue dataset. The annotated region shows one of the four tissue types. We extract and classify patches of 32 by 32 pixels.

We used convolutional RBMs with the combined learning objective to learn features from this lung tissue data.

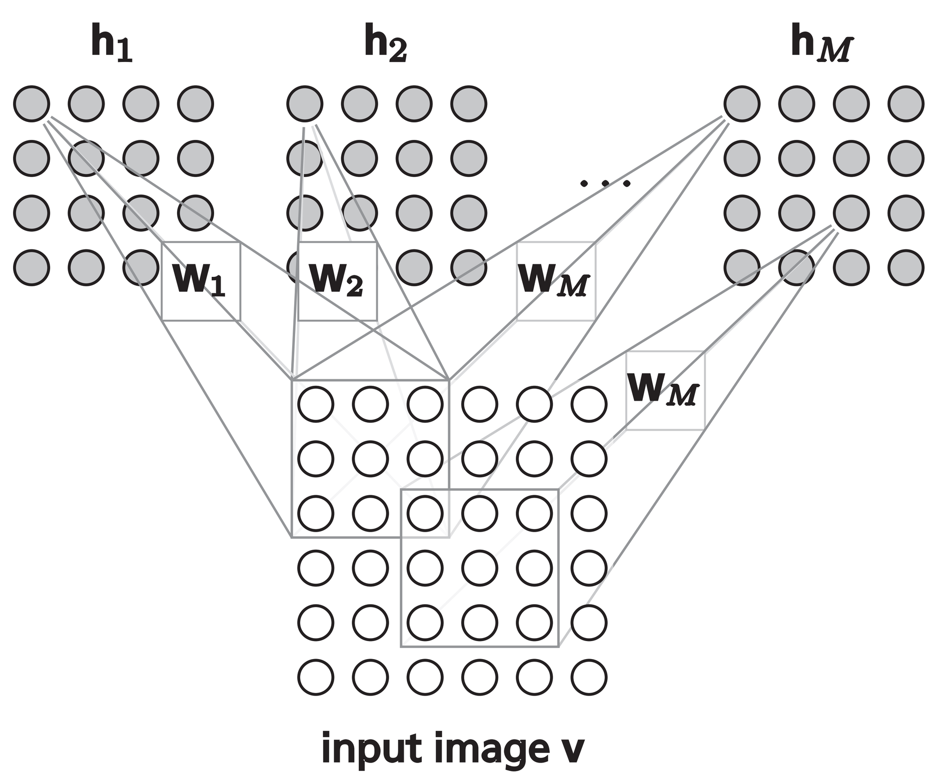

A convolutional RBM uses shared connection weights to learn a set of filters (W_i) that connect the input image to the hidden nodes, organised in feature maps (h_i).

Results

In many of our experiments, we found that a combination of generative and discriminative learning objectives gave the best results.

We trained SVMs on the features calculated by the RBMs and these SVMs usually had a higher classification accuracy than the RBMs.

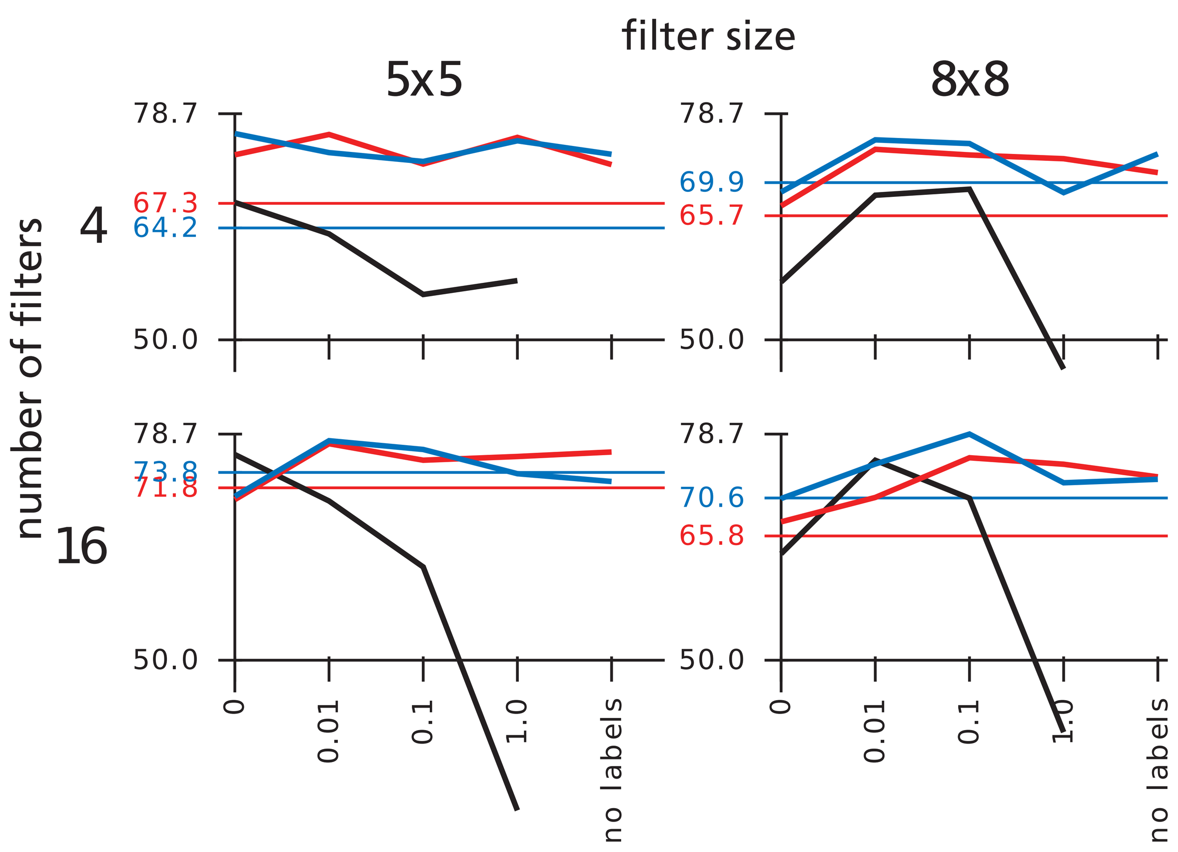

Classification accuracy of the RBM (black), and of linear (blue) and RBF (red) SVMs trained on RBM-learned features. Horizontal lines show the performance of SVM with random filters. The x-axis shows the proportion of generative learning (the β in the learning objective).

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

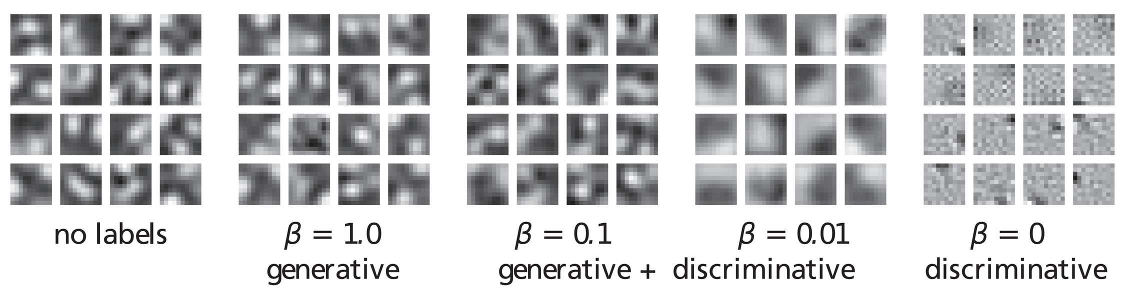

Example filters learned from the lung tissue data.

Next: other learning objectives

We want to investigate learning objectives for other applications, such as transfer learning. In medical image analysis, transfer learning could help to combine training images from different types of scanners. We hope that an RBM with a transfer learning objective might be able to learn features that improve the classification accuracy in these scenarios, because it might learn features that work for both data sources.

Reference

Hugo Larochelle et al. (2012). Learning Algorithms for the Classification Restricted Boltzmann Machine. Journal of Machine Learning Research.

Biomedical Imaging Group Rotterdam, the Netherlands

www.bigr.nl · g.vantulder@erasmusmc.nl

Summer school on deep learning for image analysis
August 2014, Denmark

Erasmus MC
University Medical Center Rotterdam