

Course 02457 Non-Linear Signal Processing, Exercise 7

This exercise is based on in C. M. Bishop: *Neural Networks for Pattern Recognition*, section 6.5–6.11. Print and comment on the figures produced by the software as outlined below at the **Checkpoints**. You are going to use the ‘Neural Classification’ Matlab toolbox programmed at Digital Signal Processing section IMM DTU

Pattern recognition with neural networks

The network should output the posterior probabilities. This is achieved using the transformation from the neural network outputs ϕ to the probabilities y by the normalized exponential or so-called *softmax* function (Bishop, pp. 215, 238+):

$$y_k = \frac{\exp(\phi_k)}{\sum_j^c \exp(\phi_j)}, \quad (1)$$

This function has a built-in redundancy: If there are c classes and we use c outputs we can calculate the output of the c 'th unit from the sum of the others. A way to overcome this redundancy is via a *modified softmax* which only has $c - 1$ outputs from the neural network (before the softmax). The c 'th output is then calculated from the others:

$$y_k = \frac{\exp(\phi_k)}{\sum_j^{c-1} \exp(\phi_j) + 1}, \quad k = 1, 2, \dots, c - 1 \quad (2)$$

$$y_c = 1 - \sum_k^{c-1} y_k \quad (3)$$

Pima indian data set

This is a data set where the task is to classify a population of women according to the risk of diabetes (binary classification). There are 7 input variables, 200 training examples and 332 test examples. 68 (34%) in the training set and 109 (32.82%) in the test set have been diagnosed with diabetes. In Brian Ripley's textbook *Pattern Recognition and Neural Networks* he states that his best method obtains about 20% misclassification on this data set so this is what we can use for reference. The input variables are:

1. Number of pregnancies
2. Plasma glucose concentration
3. Diastolic blood pressure
4. Triceps skin fold thickness
5. Body mass index (weight/height²)
6. Diabetes pedigree function
7. Age

The target output is 0 for examples having diabetes, and 1 for healthy subjects. The variables are not directly feed into the neural network. Instead the mean and standard deviation is calculated on the training set and these values are then used to normalize both the training and the test set, by first subtracting the mean and the dividing by the standard deviation.

It is not possible to view all 7 dimension in one plot. A way to plot the data is to select to variables and plot them against each other. An other way is to use principal component decomposition on the data and plot one principal component against another.

Checkpoint 7.1

Use the program `main7a.m` to optimize a neural network and plot the decision boundary and conditional probabilities.

Run the network with different weight decays and comment on the difference in the decision surfaces. Create a flow chart of `main7a.m` and its functions.

Entropic cost function

One could use the squared error function (Bishop, p. 196) to optimize the neural network for classification:

$$E_{\text{square}} = 1/2 \sum_n^N \sum_k^c (y_k^n - t_k^n)^2 \quad (4)$$

The square error costfunction is linked to ‘normal noise’ hypothesis. However, as y is suppose to model a probability and $\mathbf{t} = (0, 1, 0, \dots, 0)$ is a position coded target vector, the Gaussian error is inappropriate. We can write the conditional likelihood function:

$$p(\mathbf{t}^n | \mathbf{x}^n) = \prod_k^c [y_k^n]^{t_k^n} \quad (5)$$

This will give an more appropriate *cross-entropy* costfunction (See Bishop, p. 237). The negative log-likelihood of this is for an entire data set of independent samples:

$$E_{\text{entropy}} = - \sum_n^N \sum_k^c t_k^n \ln y_k^n \quad (6)$$

Checkpoint 7.2

The program `main7b.m` calculates the costfunction value associate with one example for a two-class problem. It sweeps the output of the neural network ϕ_1 from -10 to 10.

What was the target t_1 ?, how much do the two costfunctions penalize a “wrong” decision.

Posterior probabilities

The posterior probability that the neural network outputs can be written using Bayes theorem (here with a one-dimensional input x):

$$P(\mathcal{C}_k | x) = \frac{p(x | \mathcal{C}_k) p(\mathcal{C}_k)}{p(x)} \quad (7)$$

Checkpoint 7.3

Use the program `main7c.m` to plot posterior probabilities. The data is very simple: two classes each Gaussian distributed with $\sigma_{\mathcal{A}}^2 = \sigma_{\mathcal{B}}^2 = 1$, $\mu_{\mathcal{A}} = -2$ and $\mu_{\mathcal{B}} = 2$, thus $p(x|\mathcal{C}_{\mathcal{A}}) \sim \mathcal{N}(-2, 1)$. There are 100 samples in \mathcal{A} and 12 samples in \mathcal{B} , initially, change this to larger values (e.g. 500 and 60, to keep the prior probabilities the same) and observe how the posterior approaches the true posterior.

Pruning

Checkpoint 7.4

The program `main7d.m` will train and prune a neural network on the Pima indian set. This will take some time (and you might need to run it several times). The best net is selected as the net with the lowest costfunction value on the test set. Note, not all inputs are used, is the number of pregnancies important for the diagnosis?

The target output is 0 for examples having diabetes, and 1 for healthy subjects. What does the neural network say about the relation between age and diabetes (positive/negative correlation)?

DTU, 1999 Finn Årup Nielsen