MTAT.03.227 Machine Learning
Practice session 12

December 8 – 10, 2019

Exercise 1. Guessing coin flip results

Consider the following game:

- Your friend has 5 loaded coins which look identical. Actually, these 5 coins are different and come out heads with probabilities $0.00, 0.25, 0.50, 0.75, 1.00$, respectively. Your friend chooses uniformly randomly one of these coins and flips this coin. You can see the result (heads or tails) and have to guess what would be the result of the next flip with the same coin (still not knowing which of the 5 coins it is).

You play the above game with your friend 1000 times, each time with the same 5 loaded coins. Let us denote the actual results of the first observed flip by $y_1, \ldots, y_{1000} \in \{0, 1\}$ and the second flip that needs to be predicted by $y'_1, \ldots, y'_{1000} \in \{0, 1\}$, where 1 denotes heads and 0 denotes tails. You can choose to guess any real values in the range $[0, 1]$, let us denote these by $\hat{y}_1, \ldots, \hat{y}_{1000}$. At the end of 1000 games your guesses will be evaluated by mean squared error:

$$\mathrm{MSE} = \frac{1}{1000} \sum_{i=1}^{1000} (\hat{y}_i - y'_i)^2$$

You win the game if $\mathrm{MSE} < 0.1$ and otherwise your friend wins.
Your task in this exercise is to find the optimal strategy for yourself.

(a) Consider the strategy of always predicting the value of the first flip, that is $\hat{y}_i = y_i$. Calculate the expected error (MSE) when using this strategy.

(b) Consider the strategy of always predicting 0.5, that is $\hat{y}_i = 0.5$. Calculate the expected MSE.

(c) Now come up with a strategy that you hope would be better than the above. For this you need to decide what value $\hat{y}_i$ to predict if $y_i = 1$ and what to predict if $y_i = 0$. Calculate the expected MSE.

(d) Consider a modified game where instead of 5 coins your friend has coins with all possible loads and draws uniformly randomly one of those. That is, your friend draws a value p from the continuous uniform distribution over the interval $[0, 1]$ and then chooses the coin which has probability p to come up heads. Find the optimal strategy (no need to calculate the expected MSE).

(e) Consider the following modification to the game from subtask (d). In each of the 1000 sub-games, instead of making the prediction based on one coin flip you can now observe k consecutive flips of the same coin and need to predict the following one. Find the optimal strategy (no need to calculate the expected MSE).

Exercise 2. Bayesian linear regression. See a separate notebook on the course home page.