

Machine Learning and Data Mining
Summer 2015
Exercise Sheet 6

Presentation of Solutions to the Exercise Sheet on the 03.06.2015

Exercise 6-1 Maximum Likelihood Estimator

Suppose when soccer players train penalty kicks, each training session ends after their first miss, since they are demotivated. Let the probability of a miss be $p \in [0,1]$. Then, the probability for exactly x_i hits ($i \in \{1, \dots, N\}$) before the first miss can be modeled using the geometric distribution:

$$P(x_i) = p \cdot (1 - p)^{x_i}$$

- (a) Following a frequentist approach, determine the maximum likelihood estimator p^{ML} for an i.i.d. (independent identically distributed) population of N soccer players, which hit $x_i \in \{0, 1, \dots, \infty\}$ times before their first miss.
- (b) Consider the following dataset $\mathbf{X} = \begin{bmatrix} 7 & 2 \end{bmatrix}$ for $N = 2$ training sessions. Compute the probability of a miss.

Exercise 6-2 Frequentist versus Bayesian Statistics

Consider this – rather pathological – example to illustrate the difference between frequentist and bayesian statistics: Alice and Bob play a game in which the first person to get 6 points wins. The points are scored in the following way: A referee is standing at a pool table Alice and Bob cannot see. Before the game begins, the referee rolls a ball onto the table coming to rest at a random position. Each point scored is decided by the referee rolling another ball. If the ball comes to rest left of (the middle of) the initial ball, Alice scores, if it comes to rest right, Bob scores. The players know nothing but who scored a point. If the portion left of the initial ball is denoted as p , it is obvious that the probability of Alice scoring a point is p .

Now, consider the following situation within the game: Alice has 5 points, Bob has 3. Let us investigate the probability of Bob winning.

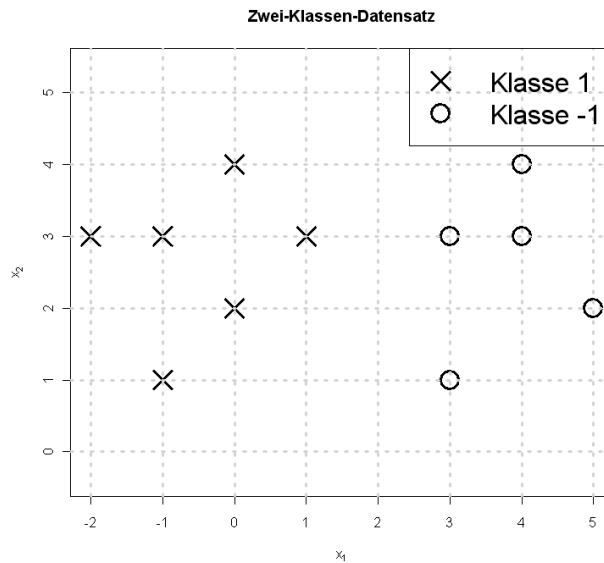
- (a) Assume that the initial ball came to rest such that $p = 2/3$. What is the probability that Bob wins?
- (b) Unfortunately, we do not know p – we only have some data we can try to estimate it from. Follow a *frequentist* approach: compute the maximum likelihood estimator for p and the probability of Bob winning.
- (c) Now, let us follow a *bayesian* approach: We know that p is only dependent on the position of the initial ball which is uniformly distributed on the table, i.e., $unif[0,1]$. Note that we compute the expected probability of Bob winning, as p itself is now drawn from a distribution. *Hint*: You will need the beta function:

$$B(x,y) = \int_0^1 t^{x-1}(1-t)^{y-1} dt = \frac{(x-1)!(y-1)!}{(x+y-1)!}$$

Optionally, you may implement the experiment and empirically compute the solution.

Exercise 6-3 Optimal Separating Hyperplane

Consider the following dataset consisting of points $\begin{pmatrix} x_1 \\ x_2 \end{pmatrix}$ in \mathbb{R}^2 . Using a hyperplane, points marked by \times are to be mapped onto ≥ 1 , points marked by \circ are to be mapped onto ≤ -1 .



- (a) Find the support vectors.
- (b) Determine the equation of the optimal separating hyperplane $h = \mathbf{x}^T \mathbf{w}$ and draw it within the figure.
- (c) Compute the margin \mathcal{C} .

Exercise 6-4 Optimal Separating Hyperplane

Determine the optimal separating hyperplane of the following dataset, partitioned into two classes A and B :

$$A = \left\{ p_1 = \begin{pmatrix} 2 \\ 4 \end{pmatrix}, p_2 = \begin{pmatrix} 3 \\ -1 \end{pmatrix}, p_3 = \begin{pmatrix} 1 \\ 0.5 \end{pmatrix}, p_4 = \begin{pmatrix} 2.5 \\ 3 \end{pmatrix}, p_5 = \begin{pmatrix} 2 \\ 2 \end{pmatrix} \right\},$$

$$B = \left\{ p_6 = \begin{pmatrix} 0.5 \\ 1.5 \end{pmatrix}, p_7 = \begin{pmatrix} -1 \\ 3 \end{pmatrix}, p_4 = \begin{pmatrix} 0 \\ 0.5 \end{pmatrix} \right\}$$

Instances of class A shall be labeled with 1, instances of class B with -1 .

Visualize the result and name the support vectors. How wide is the margin?