Machine Learning	Prof. Dr. Bernt Schiele
	Dr. Paul Swoboda & Dr. Gerard Pons-Moll
Exercise Sheet 2 - 29.10.2018	due: 5.11.2018, 14:15

Exercise 4 - Bayes Error

We have a binary classification problem, $\mathcal{Y} = \{-1, 1\}$, with the following distribution on $\mathcal{X} = [0, 1]$,

$$\mathrm{P}(Y = 1|X = x) = \begin{cases} 0.1, & \text{if } 0 \le x \le \frac{1}{8}, \\ 0.9, & \text{if } \frac{1}{8} \le x \le \frac{7}{8}, \\ 0.1, & \text{if } \frac{7}{8} \le x \le 1, \end{cases}$$

a. **(1 Points)** What is the Bayes optimal error of this problem?

b. **(1 Points)** Determine the parameter(s) (w^*, b^*) and the error of the classifier(s) $f_{(w^*,b^*)}$,

$$f_{(w,b)} = \mathrm{sign}(wx + b), \qquad w, b \in \mathbb{R}.$$

with the smallest error.

Hint:

- You don't need to give a derivation in b)- just write down the optimal parameters and the corresponding error. If there is more than one optimal set of parameters, then provide all possible optimal parameters.

Exercise 5 - Loss functions and Bayes optimal functions

a. **(3 Points)** Let $\mathcal{Y} = \{-1, 1\}$ (binary classification). Show that the Bayes optimal function, $f^*(x) = \underset{c \in \mathbb{R}}{\arg\min}\, \mathbb{E}[L(Y, c)|X = x]$, for the least squares loss, $L(y, f(x)) = (y - f(x))^2$, is $f^*(x) = \mathbb{E}[Y|X = x]$ and deduce that the least squares loss is classification calibrated.

b. **(2 Points)** Let $\mathcal{Y} = \mathbb{R}_+ = \{x \in \mathbb{R}|x \ge 0\}$ (regression with output on the positive part of $\mathbb{R}$) and suppose that $\mathbb{E}[Y|X = x] > 0$. Show that the Bayes optimal function, $f^*(x) = \underset{c \in \mathbb{R}}{\arg\min}\, \mathbb{E}[L(Y, c)|X = x]$ for the loss function $L(y, f(x)) = \log(f(x)) + \frac{y}{f(x)}$, is given by $f^*(x) = \mathbb{E}[Y|X = x]$. Discuss the properties of this loss function compared to least square loss. For what kind of noise model do you think is this loss function useful (note that the target space is the set of non-negative reals) ?

Exercise 6 - Maximum Likelihood and Maximum A Posteriori Estimation

We have as likelihood function,

$$p(x\,|\,\theta) = \frac{1}{\sqrt{2\pi}}\, e^{-\frac{(x-\theta)^2}{2}},$$

which is a Gaussian density with mean θ and unit variance.

a. **(3 Points)** Derive the maximum-likelihood estimate of θ given n independent samples $x_1, \ldots, x_n$.

b. **(3 Points)** We have now an additional prior for the value of θ given as

$$p(\theta) = \begin{cases} \gamma\, e^{-\gamma\theta}, & \theta \geq 0, \\ 0, & \theta < 0. \end{cases},$$

for some fixed value of $\gamma > 0$. Derive the MAP estimator of θ.

c. **(1 Point)** Is the prior $p(\theta)$ reasonable ? Suppose that the data is sampled from a Gaussian distribution with unit variance. Does the MAP estimator converge to the true mean parameter of the Gaussian as $n \to \infty$?

Submission instructions

- We accept both handwritten and electronic submissions. So you can choose what is more convenient for you. In any case, you should specify full names and immatriculation IDs of all team members. Obviously, programming tasks you can submit only electronically.
- Handwritten submissions should be submitted in the lecture hall of Monday's lecture (before the lecture starts).
- Electronic submissions should be emailed to the corresponding tutor:

  a. Apratim Bhattacharyya (Wednesday 8-10): abhattac@mpi-inf.mpg.de

  b. Maksym Andriushchenko (Thursday 8-10): s8mmandr@stud.uni-saarland.de

  c. Max Losch (Friday 16-18): mlosch@mpi-inf.mpg.de

  If not all 3 students belong to the same tutorial group, then you should email your submission to **only** one tutor (e.g. to the tutor of the first author of your homework), so please do not put other tutors in copy of the email.

  The email subject must have the following form: "[ML18/19 Exercise] Sheet X", where X is the number of the current exercise sheet. Then please specify in the email full names and immatriculation IDs of all team members. Then please attach all your files as a single zip archive, which consists of your immatriculation IDs, e.g. "2561234_2561235_2561236.zip".
- Reminder: you should submit in groups of 3. Otherwise, we will later on merge the groups smaller than 3 students.