SGN-41007 Pattern Recognition and Machine Learning
Exam 9.10.2019
Heikki Huttunen

- ▷ Use of calculator is allowed.
- ▷ Use of other materials is not allowed.
- ▷ The exam questions need not be returned after the exam.
- ▷ You may answer in English or Finnish.

1. Are the following statements true or false? No need to justify your answer, just T or F. Correct answer: 1 pts, wrong answer: $-\frac{1}{2}$ pts, no answer 0 pts.

   (a) Maximum likelihood estimators are unbiased.

   (b) The Receiver Operating Characteristics curve plots the probability of detection versus the probability of false alarm for all thresholds.

   (c) Least squares estimator minimizes the squared distance between the data and the model.

   (d) The number of support vectors of a support vector machine equals the total number of samples.

   (e) The LDA maximizes the variance of samples in each classes.

   (f) Cross-validation is used for model accuracy evaluation.

2. The *Poisson distribution* is a discrete probability distribution that expresses the probability of a number of events $x \geq 0$ occurring in a fixed period of time:

$$p(x;\lambda) = \frac{e^{-\lambda}\lambda^x}{x!}$$

   We measure N samples: $x_0, x_1, \ldots, x_{N-1}$ and assume they are Poisson distributed and independent of each other.

   (a) Compute the probability $p(\mathbf{x};\lambda)$ of observing the samples $\mathbf{x} = (x_0, x_1, \ldots, x_{N-1})$. (1p)

   (b) Compute the natural logarithm of p, *i.e.*, $\log p(\mathbf{x};\lambda)$. (1p)

   (c) Differentiate the result with respect to λ. (2p)

   (d) Find the maximum of the function, *i.e.*, the value where $\frac{\partial}{\partial\lambda}\log p(\mathbf{x};\lambda) = 0$. (2p)

	Prediction	True label
Sample 1	0.8	1
Sample 2	0.5	1
Sample 3	0.6	0
Sample 4	0.1	0

Table 1: Results on test data for question 5a.

3. Two measurements $x(n)$ and $y(n)$ depend on each other in a linear manner, and there are the following measurements available:

n	0	1	2
$x(n)$	7	9	4
$y(n)$	12	15	4

We want to model the relationship between the two variables using the model:

$$y(n) = ax(n) + b.$$

Find the L_2-regularized least squares estimates $\hat{a}$ and $\hat{b}$ that minimize the squared error using penalty $\lambda = 10$.[1]

4. (6 pts) Consider the Keras model defined in Listing 1. Inputs are 64×64 color images from 10 categories.

   (a) Draw a diagram of the network.

   (b) Compute the number of parameters for each layer, and their total number over all layers.

5. (a) (4p) A random forest classifier is trained on training data set and the `predict_proba` method is applied on the test data of Table 1. Draw the receiver operating characteristic curve. What is the Area Under Curve (AUC) score?

   (b) (2p) A binary classifier is trained with 1 million samples from two classes. The AUC of the classifier on test data with another 1 million samples, is 0.768. We choose one sample from the positive class at random and another three samples from the negative class at random. What is the probability that the sample from the positive class has highest score of the four samples [hint: study the literature on the last page]?

[1] Alternatively, the unregularized solution will give you max. 4 points.

Listing 1: A CNN model defined in Keras

```
model = Sequential()

w, h = 3, 3
sh = (64, 64, 3)

model.add(Convolution2D(32, w, h, input_shape=sh, border_mode='same'))
model.add(MaxPooling2D(pool_size=(2, 2)))
model.add(Activation('relu'))

model.add(Convolution2D(32 w, h, border_mode='same'))
model.add(MaxPooling2D(pool_size=(2, 2)))
model.add(Activation('relu'))

model.add(Convolution2D(48, w, h, border_mode='same'))
model.add(MaxPooling2D(pool_size=(2, 2)))
model.add(Activation('relu'))

model.add(Convolution2D(48, w, h, border_mode='same'))
model.add(MaxPooling2D(pool_size=(2, 2)))
model.add(Activation('relu'))

model.add(Flatten())
model.add(Dense(128))
model.add(Activation('relu'))

model.add(Dense(10, activation = 'softmax'))
```