

SGN-41007 Pattern Recognition and Machine Learning
Exam 27.2.2018
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) Least squares estimator minimizes the squared distance between neighboring samples.
- (b) The Receiver Operating Characteristics curve plots the probability of detection versus the probability of false alarm for all thresholds.
- (c) The LDA maximizes the following score:

$$J(\mathbf{w}) = \frac{\text{Distance of class means}}{\text{Variance of classes}}$$

- (d) A neural network classifier has a linear decision boundary between classes.
 - (e) L1 and L2 regularization improves the generalization of a logistic regression classifier.
 - (f) Maxpooling returns the maximum over input channels.
2. 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	2
$y(n)$	11.6	14.8	3.5

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

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

ROC space [edit]

The contingency table can derive several evaluation "metrics" (see infobox). To draw a ROC curve, only the true positive rate (TPR) and false positive rate (FPR) are needed (as functions of some classifier parameter). The TPR defines how many correct positive results occur among all positive samples available during the test. FPR, on the other hand, defines how many incorrect positive results occur among all negative samples available during the test.

A ROC space is defined by FPR and TPR as x and y axes respectively, which depicts relative trade-offs between true positive (benefits) and false positive (costs). Since TPR is equivalent to sensitivity and FPR is equal to $1 - \text{specificity}$, the ROC graph is sometimes called the sensitivity vs ($1 - \text{specificity}$) plot. Each prediction result or instance of a confusion matrix represents one point in the ROC space.

Another complication in applying LDA and Fisher's discriminant to real data occurs when the number of measurements of each sample (i.e., the dimensionality of each data vector) exceeds the number of samples in each class.^[4] In this case, the covariance estimates do not have full rank, and so cannot be inverted. There are a number of ways to deal with this. One is to use a pseudo inverse instead of the usual matrix inverse in the above formulae. However, better numeric stability may be achieved by first projecting the problem onto the subspace spanned by Σ_b .^[12] Another strategy to deal with small sample size is to use a shrinkage estimator of the covariance matrix, which can be expressed mathematically as

$$\Sigma = (1 - \lambda)\Sigma + \lambda I$$

where I is the identity matrix, and λ is the *shrinkage intensity* or *regularisation parameter*. This leads to the framework of regularized discriminant analysis^[13] or shrinkage discriminant analysis.^[14]