Department of Mathematics and Computer Science
University of Southern Denmark, Odense

February 1, 2013
Marco Chiarandini

DM825 - Introduction to Machine Learning

Sheet 2, Spring 2013

Prepare exercises 1.11, 1.14, 1.24, 3.8 from book [B1] of the course literature

Exercise 1
Suppose that a fair-looking coin is tossed three times and lands heads each time. Show that a classical maximum likelihood estimate of the probability of landing heads would give 1, implying that all future tosses will land heads. By contrast, show that a Bayesian approach with a prior of 0.5 for the probability of heads would lead to a much less extreme conclusion on the posterior probability of observing heads.

Exercise 2
Show the derivation of the results for μ_m and $1/\sigma_m$ presented on slide 21 of today's lecture.

Exercise 3. Linear Regression and k nearest neighbor The files q2x.dat and q2y.dat contain the inputs x^i and outputs y^i for a regression problem, with one training example per row.

i. Implement the linear regression ($y = \theta^T x$) on this dataset using the normal equations (which is done in R automatically via the `lm` function) and plot on the same figure the data and the straight line resulting from your fit (in R, plot the points and then pass the fitted linear model to `abline`). Compare your result with the implementation via the sequential gradient algorithm from the past exercise sheet. (Remember to include the intercept term.)

ii. Implement locally weighted linear regression on this dataset and plot on the same figure

iii. Implement a k-nearest neighbor regression (in R install package FNN and read the documentation of `knn.reg`). Use some randomly chosen x values as test points. Plot the training and predicted points for $k = 3$. Further, show graphically the behavior of the square error as k increases from $k = 0$ to the size of the training set that you decided.