

DM825 - Introduction to Machine Learning

Sheet 1, Spring 2013

Exercise 1

Consider the application of Linear Models in classification. For example, consider the application in medical diagnosis of diabetes. The features are blood sugar measurement and amount of fluid intake per day of different patients. These features are used to classify the symptoms sets into two different diagnosis categories, that is, presence or not of diabetes.

Table 1 reproduces some past measurements with the relative category. Here +1 means that the person had diabetes, while -1 that he/she did not. (data.txt)

Implement in R the Batch and Stochastic Gradient Algorithms to learn the weights of x_1, x_2 and the *bias* in the linear model such that it outputs outputs $-1, +1$. Report the final regression coefficients and plot the points together with the linear separator.

x_1	x_2	Diabetes
1.0	1.0	1
9.4	6.4	-1
2.5	2.1	1
8.0	7.7	-1
0.5	2.2	1
7.9	8.4	-1
7.0	7.0	-1
2.8	0.8	1
1.2	3.0	1
7.8	6.1	-1

Table 1: The data for Exercise 1.