

Deadline: Fr. Mai 24, 14:00 Drop your printed or legible handwritten submissions into the boxes at Samelsonplatz, or a .pdf file via LearnWeb.

1 Model Averaging - XOR (10 + 4 points)

A. [4p] Show that voting and averaging do not always yield the same result by coming up with a (small) example.

B. [4p] Consider the XOR-Problem (Table 1). We know that a single linear classifier such as a Logistic Regression model cannot solve the XOR-problem with 100% accuracy. Explain how it can be solved by combining multiple Logistic Regression models through BAGGING.

C. [2p] Perform an experiment to verify this result by training 100 Logistic Regression models on randomly chosen subsets of the XOR-dataset of size $N = 3$. (You can use a software library such as SCIKIT-LEARN)

D. Bonus (hard) [4p] An alternative idea to parts A and B would be to train 100 Logistic Regression models on the full XOR-dataset, but with randomly chosen initial parameters (before fitting), and then averaging them. Explain why this won't work.

x_1	x_2	y
-1	-1	0
-1	+1	1
+1	-1	1
+1	+1	0

Table 1: XOR-dataset

2 Random Forests (10 points)

Train a Random Forest with 3 decision stumps using the Gini-Index as the splitting criterion. Use the following configuration:

- Stump 1: train on instances 2,3, features: x_2, x_3
- Stump 2: train on instances 1,2, features: x_1, x_2
- Stump 3: train on instances 3,4, features: x_1, x_3

Finally, make the prediction for the instances 5 and 6.

x_1	x_2	x_3	y
1	1	1	1
1	1	0	1
0	1	1	0
0	0	0	0
1	0	1	?
0	1	0	?