

Homework 1 (by December 5, 2017)

1. practice building and running tensorflow graphs
2. implement **Logistic Regression**, i.e. learn w by minimizing the *logistic* loss

$$L(w) = \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-y_i w^\top x_i))$$

using a) fixed data, and b) data being handed in via placeholders

3. try different learning rates to find one that converges faster than $\eta = 0.001$
4. (optional) create a version with analytically computed gradients, compare it speed

hand-in requirement

5. upload your code to 2a),2b) with a reasonable η to the *git* server