

Statistical methods in NLP

Evaluating and comparing classifiers



Richard Johansson

February 16, 2016

overview

overview of the second programming assignment

significance testing for the accuracy


comparing two classifiers

the second programming assignment

- ▶ in the previous assignment, you implemented a classifier
- ▶ you will now make a more thorough analysis of how well it works

learning curve

- ▶ set aside a test set
- ▶ use training sets of different sizes
- ▶ estimate the accuracy for each training set size
- ▶ plot the result



overview


overview of the second programming assignment

significance testing for the accuracy

comparing two classifiers

example

- ▶ 40 correct guesses out of 50
- ▶ is the true accuracy really different from 0.70?



overview

overview of the second programming assignment

significance testing for the accuracy

comparing two classifiers

