

One-Pass Boosting

Z. Barutcuoglu (Princeton), P.M. Long (Google), R.A. Servedio (Columbia)

Motivation. Common boosting scenario: at each stage the *best-performing* base classifier is chosen from a fixed pool of candidates.

If pool is very large (n -grams, amino acid sequences), may be too expensive to optimize at each stage.

This paper: boosting algorithms that make *one pass* over pool of base classifiers

- In one-pass setting it can pay to be picky
- For “diverse base classifiers” scenario, we get same performance guarantee as (more computationally intensive) multi-pass boosting using one-pass algorithm
- Experimental results