CS156b Project1

Letter Recognition — Boosting

Ling Li
Learning Systems Group, Caltech
November 6, 2001

Last week

- A bug in adaptive learning program fixed
- Tuned the program with `gprof`
  - most time-consuming parts: `tanh()`, feed-forward, back-propagation
  - replaced `tanh()` with two look-up tables
  - rewrote feed-forward and back-propagation: more pointer operations
  - decreased the precision of line-search (which decreases the # of feed-forward calls)

  The training is now twice fast, while the performance (training/validation errors) remains.
- Tried bagging, simple boost, and AdaBoost.M1
  - Base model: 16-50-26, normalized data and $e^{(3)}$ (W3) error function, momentum, 700 epochs.

Bagging

boostrap **agg**regat**ing**: resample with replacement, unweighted average

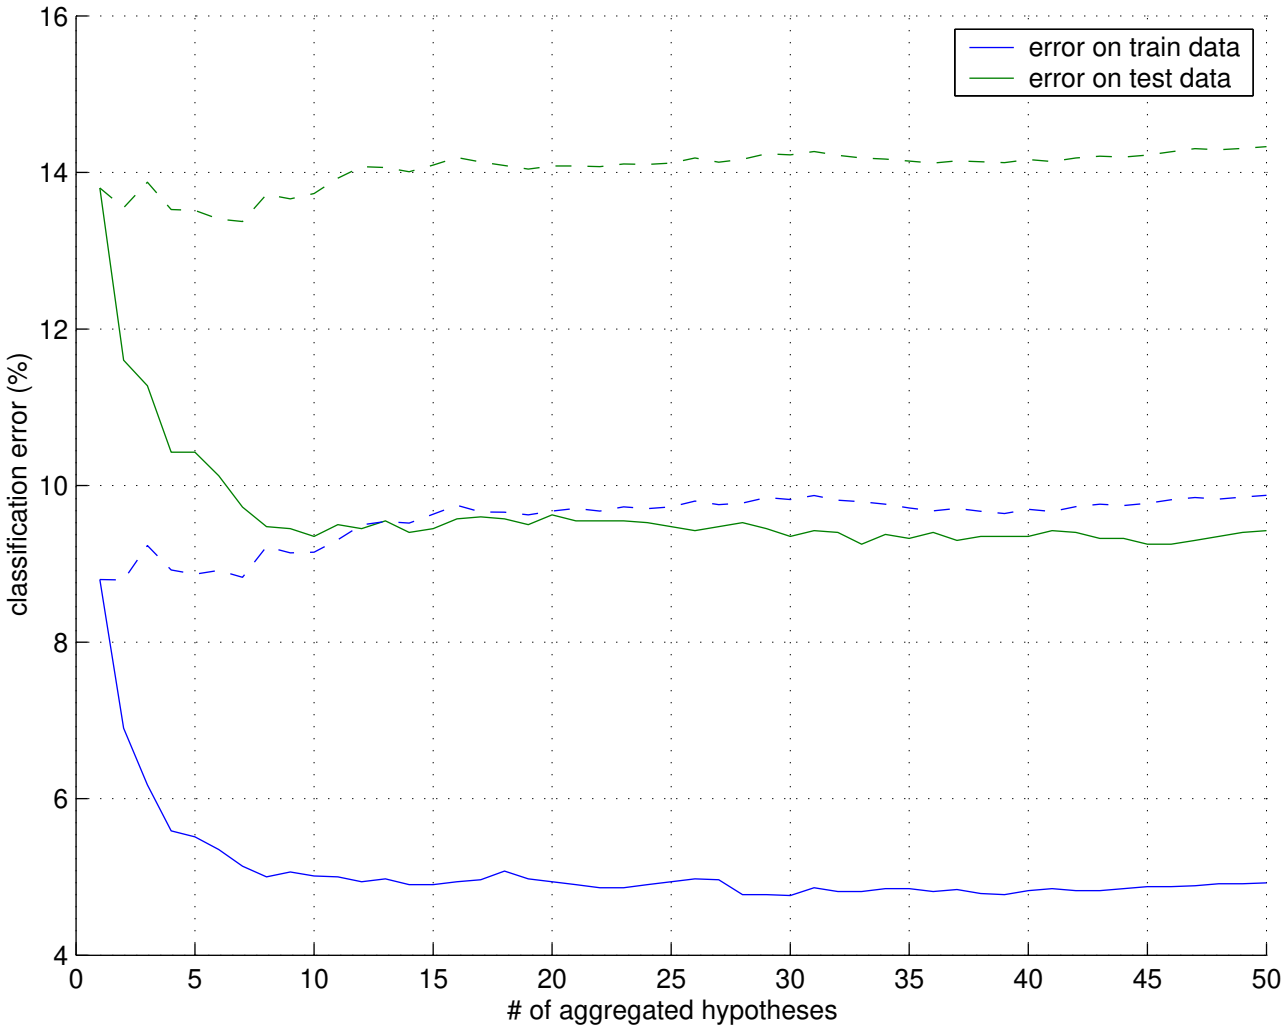

Simple boost

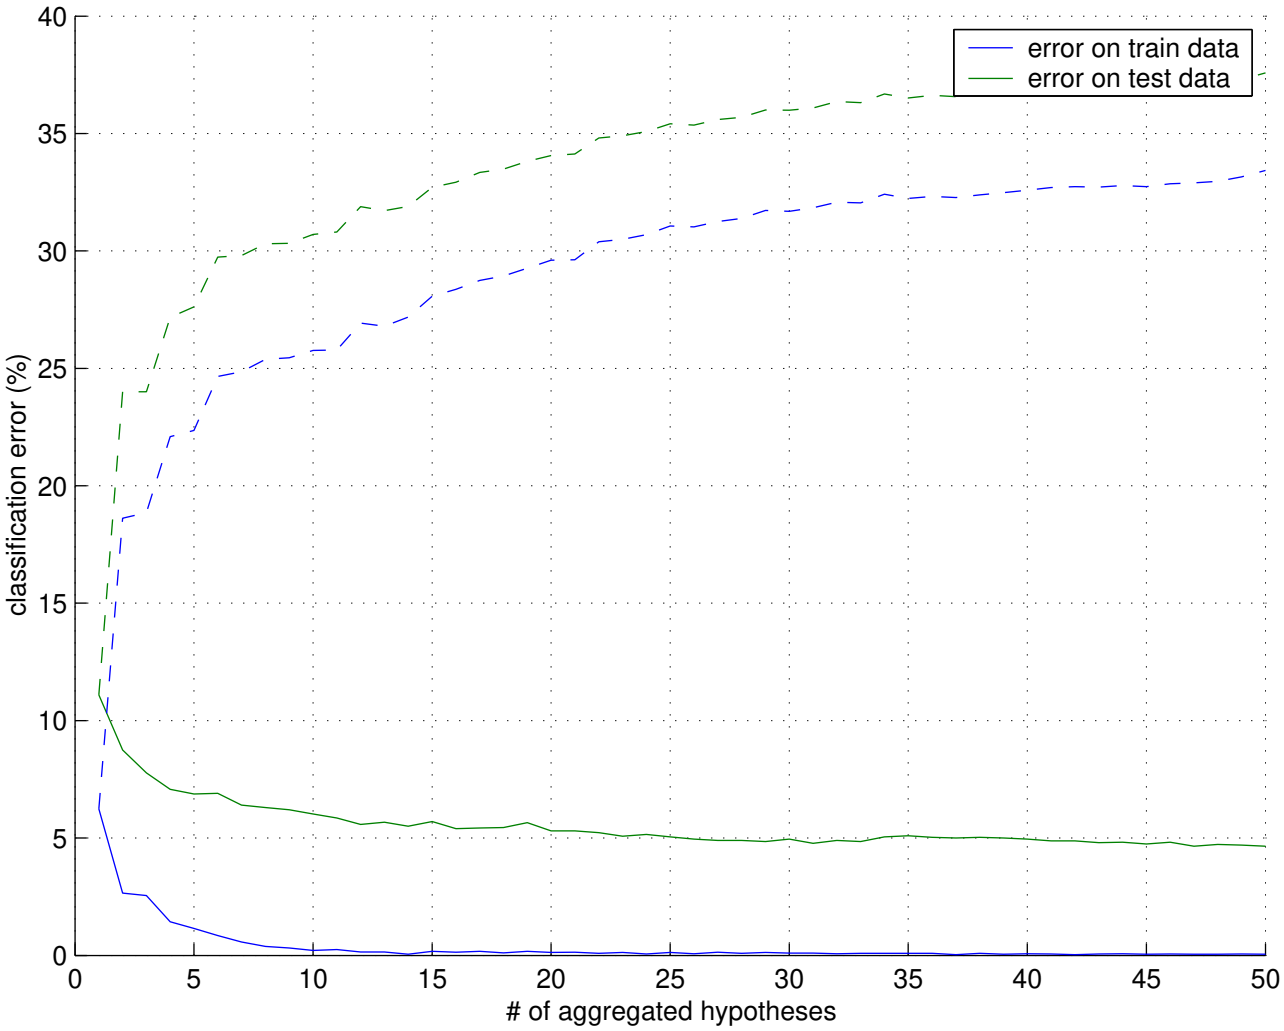

AdaBoost.M1

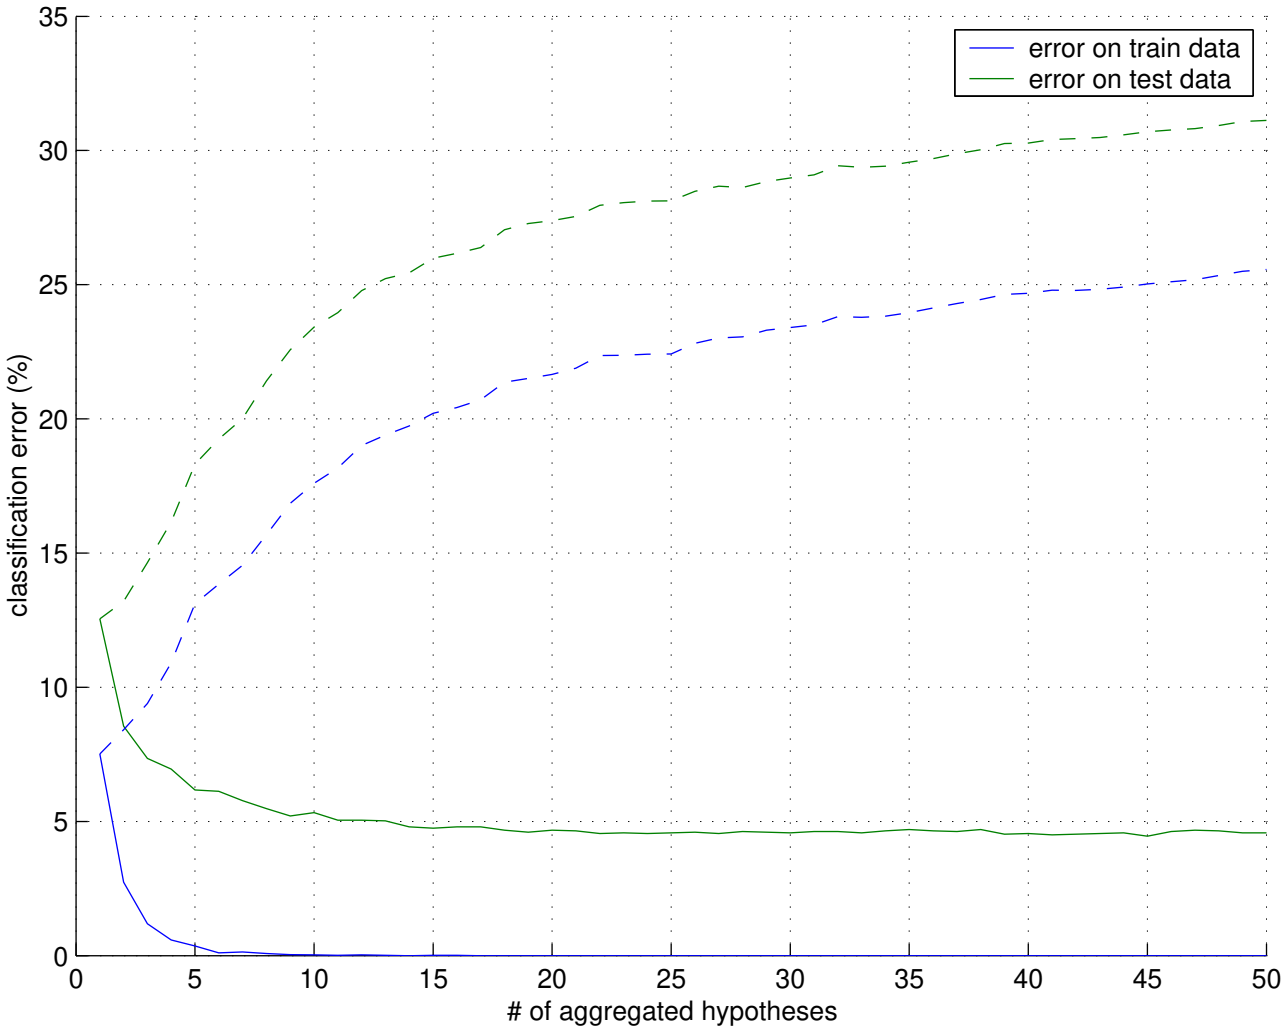
error on train data
error on test data
classification error (%)
# of aggregated hypotheses
0
5
10
15
20
25
30
35
40
45
50