


CSE 446

Ensembles


Administrative

- Grading
 - Homework 1 grades are out
 - Midterm grading in progress
- Homework 2 due today
- Homework 3 already out, start early!
- Today: model ensembles

Dataset



Decision Tree Fit



Voting (Ensemble Methods)

- Instead of learning a single classifier, **learn many weak classifiers** that are **good at different parts of the data**
- **Output class:** (Weighted) vote of each classifier
 - Classifiers that are most “sure” will vote with more conviction
 - Classifiers will be most “sure” about a particular part of the space

Simple version: BAGGing = Bootstrap AGGregation


(Breiman, 1996)

- for $t = 1, 2, \dots, T$:
 - $D_t \leftarrow$ randomly select M training instances with replacement
 - $h_t \leftarrow \text{learn}(D_t)$ [*Decision Tree, Naive Bayes, ...*]
- Now combine the h_t together with uniform voting ($w_t = 1/T$ for all t)


$D = \{ [0, 0, 1 \mid 0], [1, 1, 0 \mid 0], [0, 1, 0 \mid 1], [1, 1, 1 \mid 0] \}$

$D_t = \{ [1, 1, 0 \mid 0], [1, 1, 0 \mid 0], [1, 1, 1 \mid 0], [0, 0, 1 \mid 0] \}$


Dataset



Decision Tree Fit

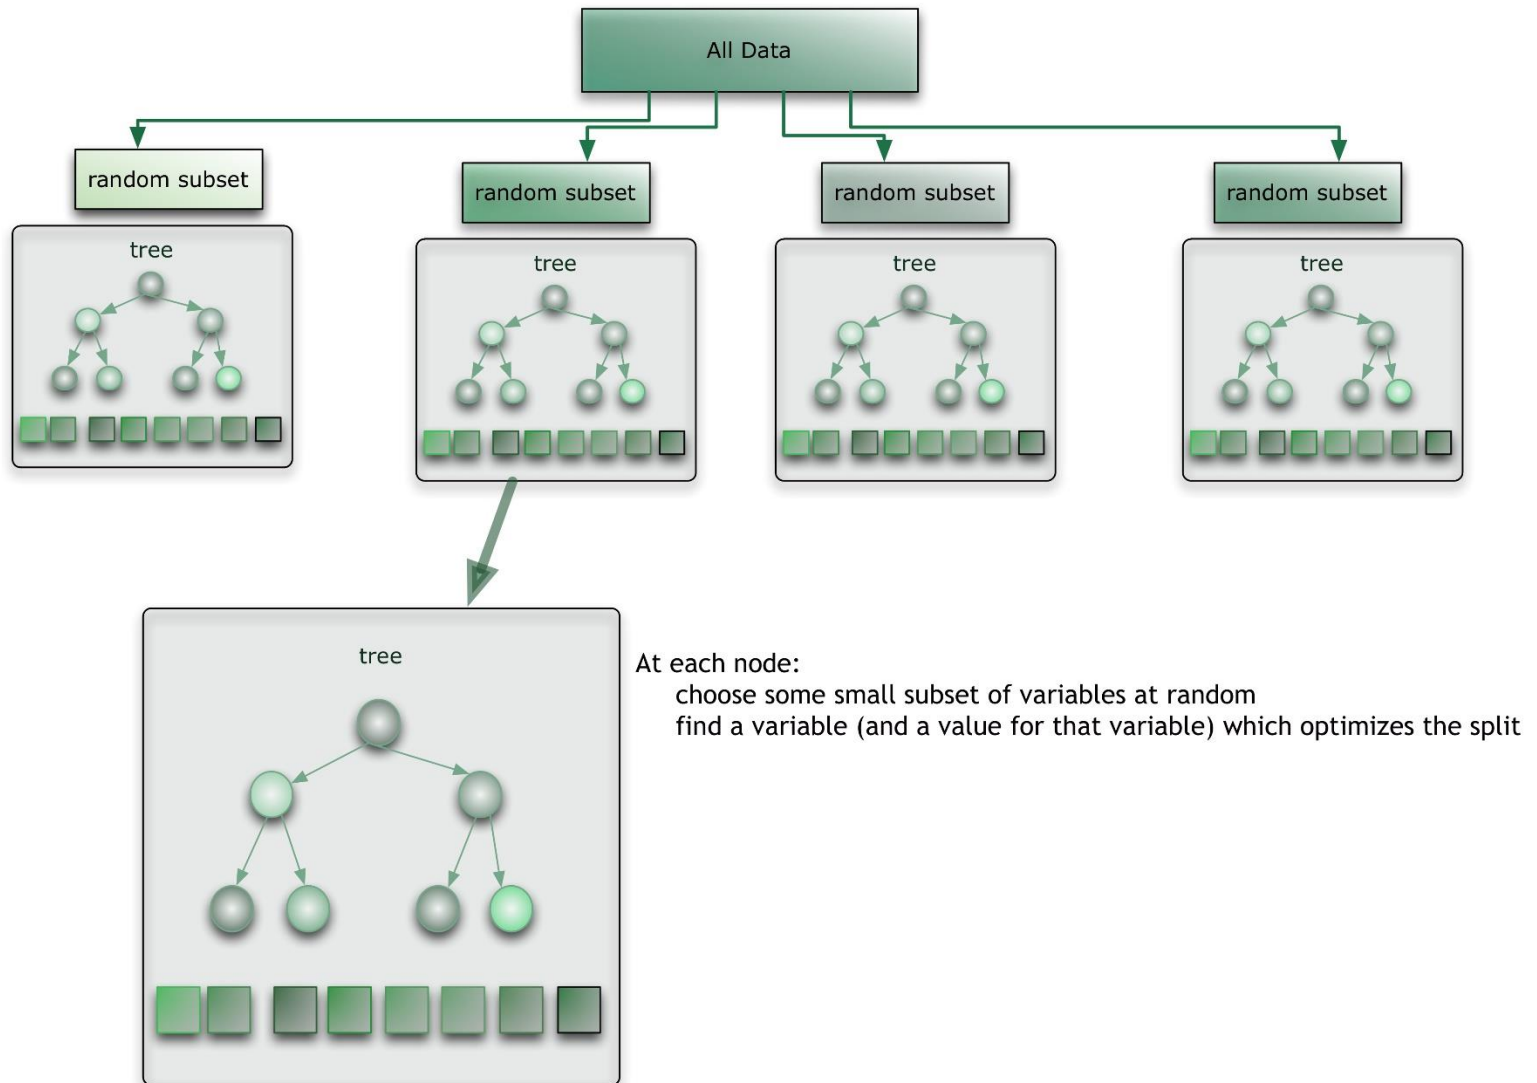


100 bagged trees



shades of blue/red indicate strength of vote for particular classification

Random Forests



Fighting the bias-variance tradeoff

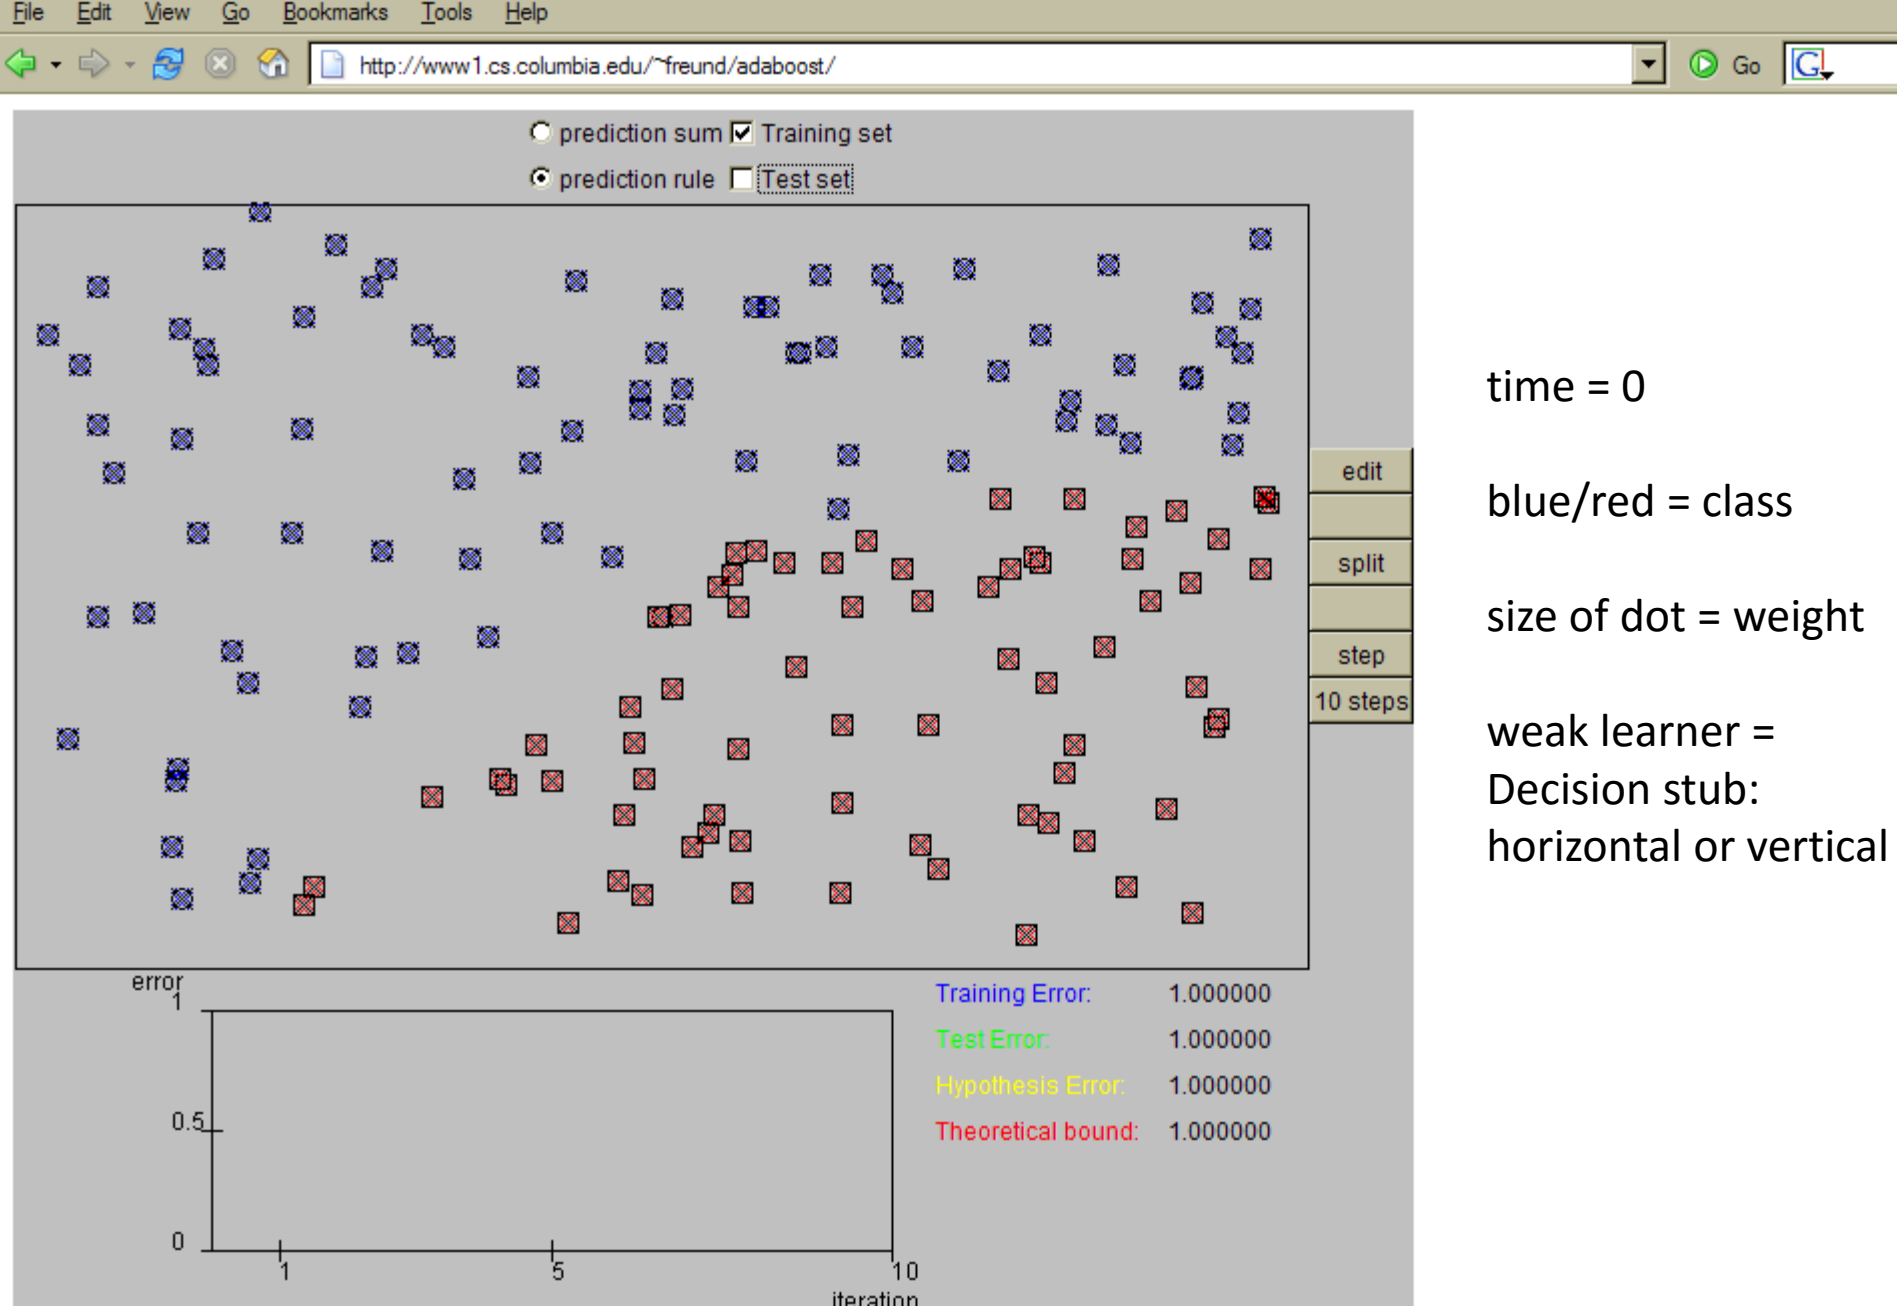
- **Simple (“weak”) learners**
 - e.g., naïve Bayes, logistic regression, decision stumps (or shallow decision trees)
 - Low variance, don’t usually overfit
- **Why not use weak learners all the time?**
 - High bias, can’t solve hard learning problems
- **Ensembles use independent weak learners (which don’t overfit as much), and put many of them together to reduce bias**

Boosting

[Schapire, 1989]

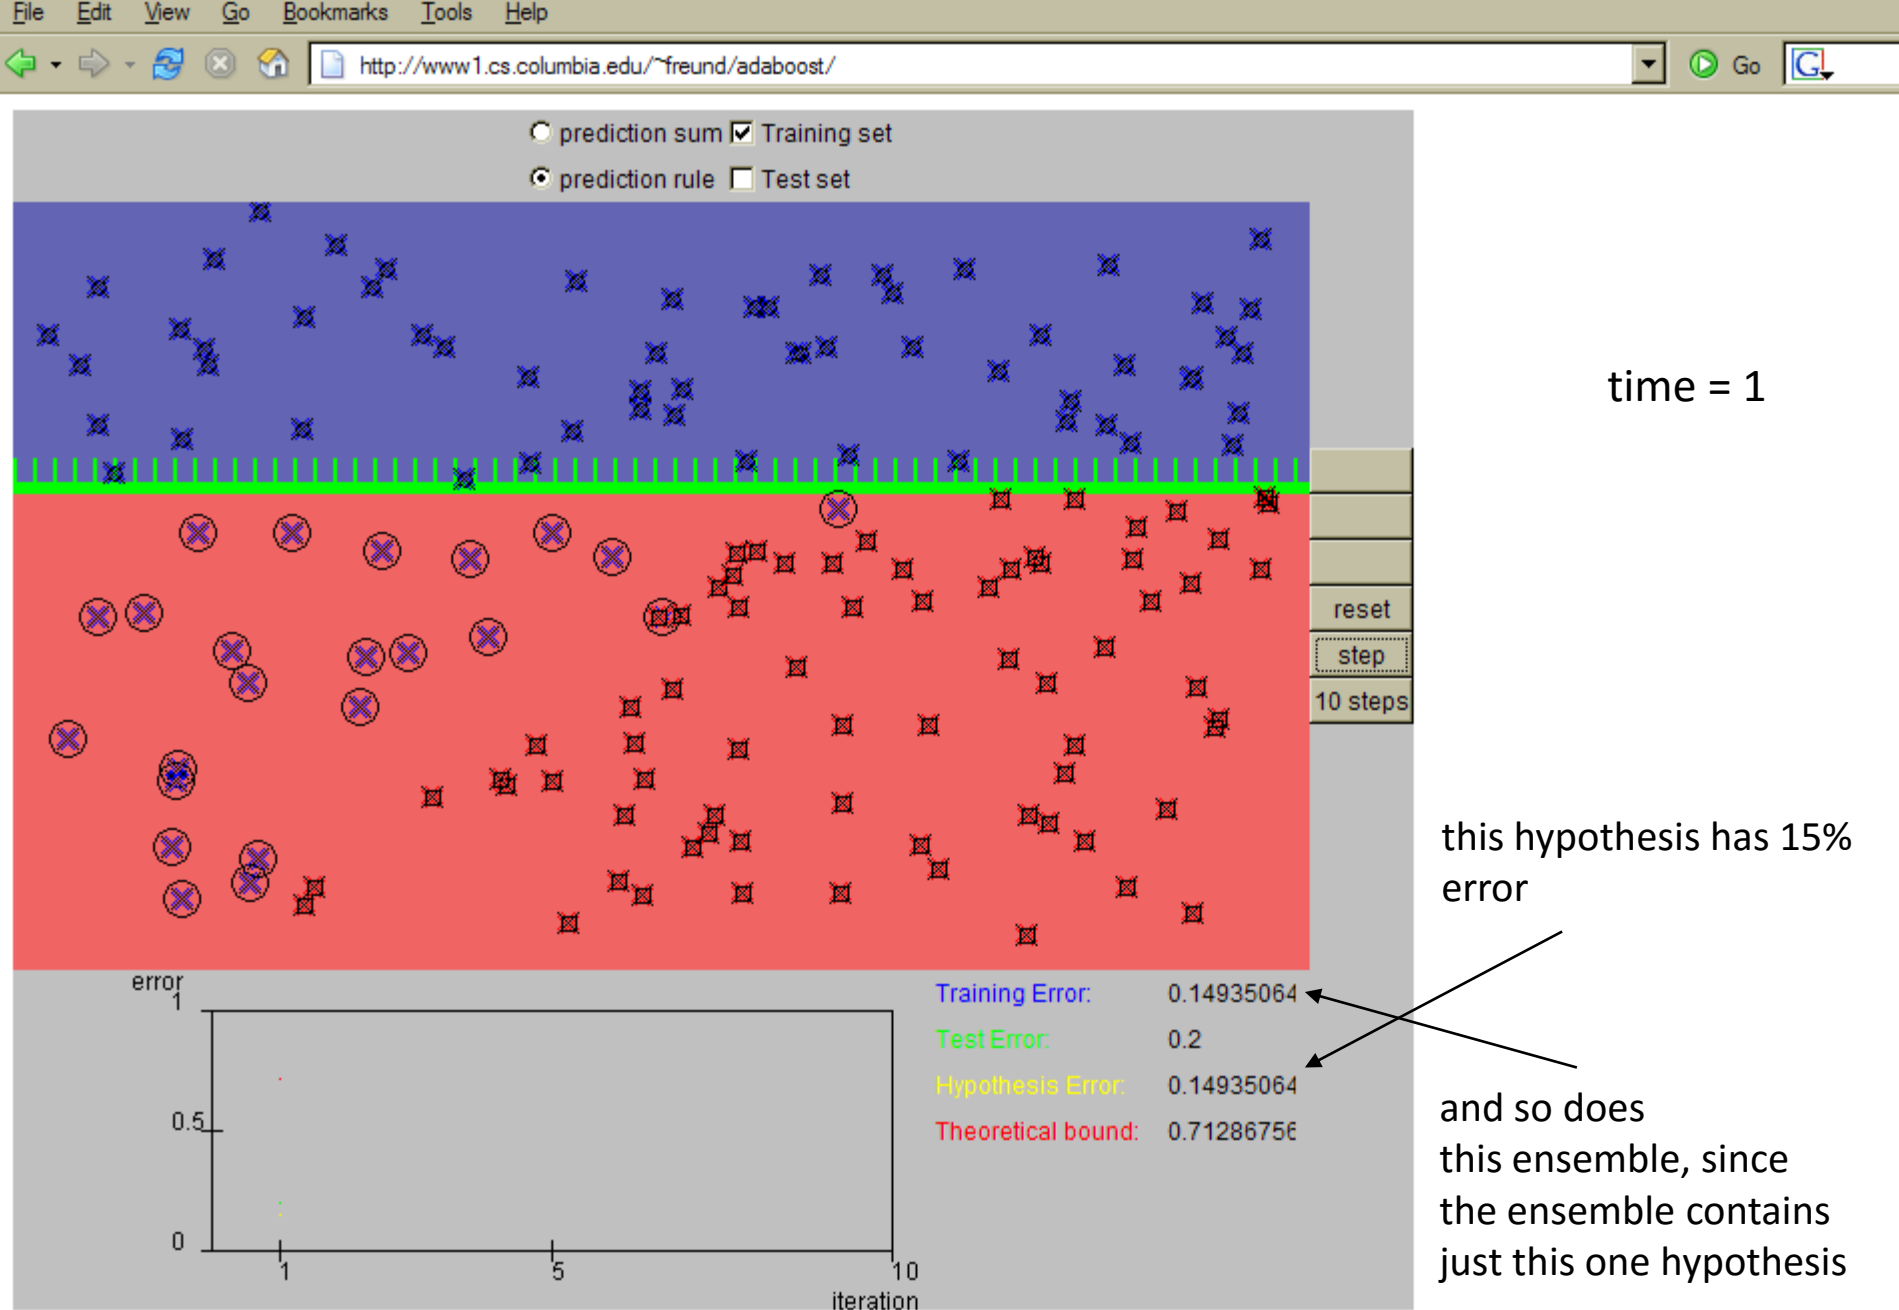
- **Idea:** if we give each weak learner a different piece of the dataset, we get a really good complex classifier from letting them vote
- Learners must be different (how was this achieved in Bagging?)
- Learners must be better than random (not too weak)
- Approach: given a weak learner, run it multiple times on (reweighted) training data, then let learned classifiers vote
- On each iteration t :
 - weight each training example by how incorrectly it was classified
 - Learn a hypothesis – h_t
 - A strength for this hypothesis – α_t
- Final classifier:

$$h(x) = \text{sign} \left(\sum_i \alpha_i h_i(x) \right)$$
- **Practically useful**
- **Theoretically interesting**




First, generate a data-set by clicking on the left and right buttons in the main window of the applet. Then, press "split" to split the data into training and test sets

Applet adaboost started



First, generate a data-set by clicking on the left and right buttons in the main window of the applet. Then, press "split" to split the data into training and test sets




time = 2

First, generate a data-set by clicking on the left and right buttons in the main window of the applet. Then, press "split" to split the data into training and test sets




Applet adaboost started



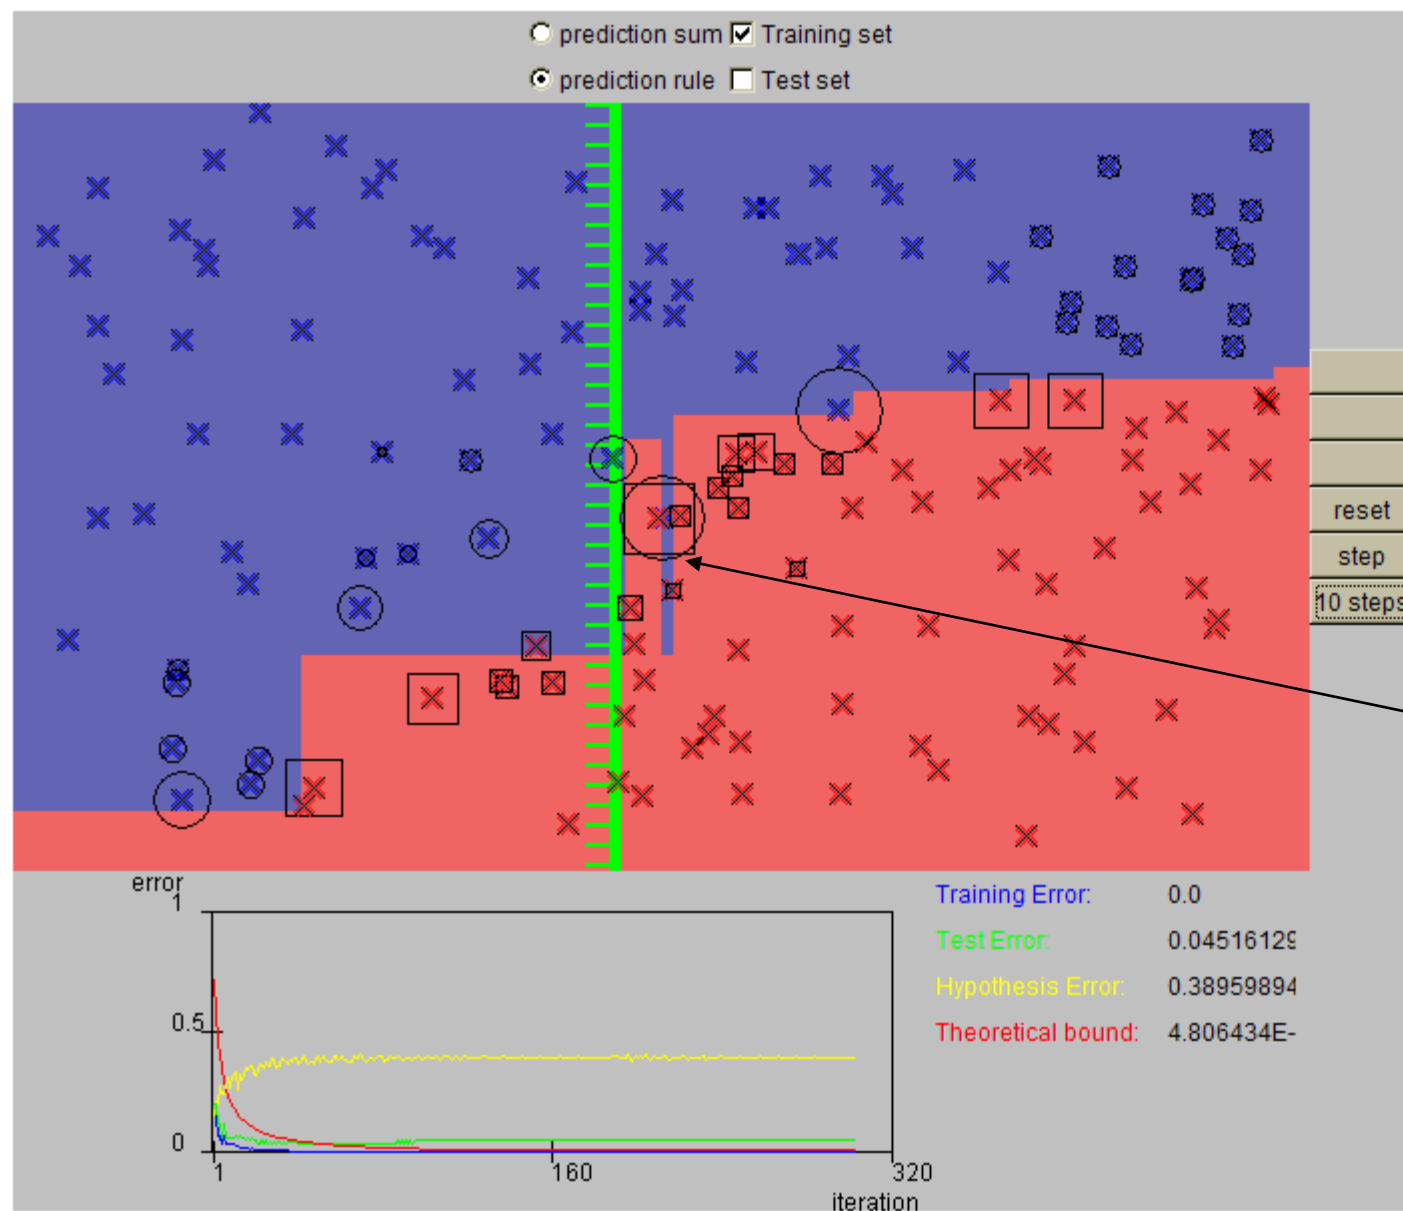
time = 13

First, generate a data-set by clicking on the left and right buttons in the main window of the applet. Then, press "split" to split the data into training and test sets



time = 100

First, generate a data-set by clicking on the left and right buttons in the main window of the applet. Then, press "split" to split the data into training and test sets



time = 300

overfitting

First, generate a data-set by clicking on the left and right buttons in the main window of the applet. Then, press "split" to split the data into training and test sets.