LEARNING ENSEMBLES

Learning model ensembles

- **Basic idea:**
  Instead of learning one model,
  Learn several and combine them
- Typically improves accuracy, often by a lot
- **Many methods:**
  - Bagging
  - Boosting
  - ECOC (error-correcting output codes)
  - Stacking
  - Etc.

Bagging

- Generate "bootstrap" replicates of training set by sampling with replacement
- Learn one model on each replicate
- Combine by uniform voting

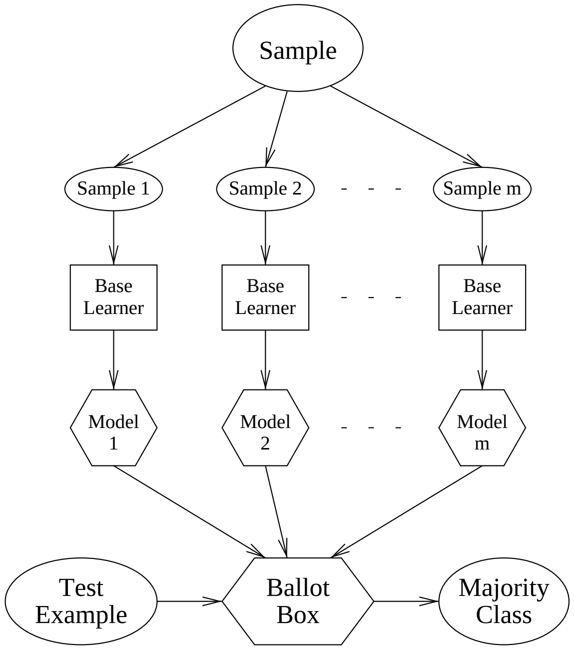

Boosting

- Maintain vector of weights for examples
- Initialize with uniform weights
- Loop:
  - Apply learner to weighted examples (or sample)
  - Increase weights of misclassified examples
- Combine models by weighted voting

Fig. 2 from (Dietterich, 1997)

Issues / Open questions

- Boosting tends to outperform bagging
- But it's more sensitive to noise
- Why do they work?
- Error = Bias + Variance
  (Bagging reduces variance; boosting also reduces bias)
- How can we recover comprehensibility?
- Can we make ensemble learning more efficient?
- Etc.