SAND2007-7463C

Why and How to exploit OOB Validation for Ensemble Size

Philip Kegelmeyer, Sandia National Labs, wpk@sandia.gov

CASIS, November 16, 2007

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energys National Nuclear Security Administration under contract DE-AC04-94AL85000.

SAND Document #2007-6608A, R&A #5256574

Machine Learning, With Ensembles

Traditional: Use 100% of training data to build a sage.

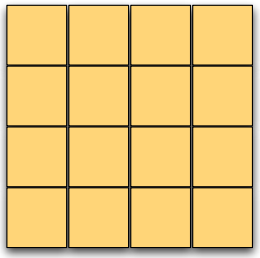

Sage sees all the data.

Ensembles: Use randomized 100% of training data to build an expert. Repeat to build many experts. Vote them.

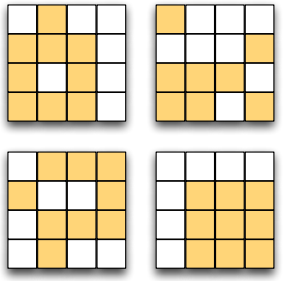

Each expert sees 2/3rds of the data.

The experts beat the sage[1]!

"Bagging" is the Formal Name for This Method

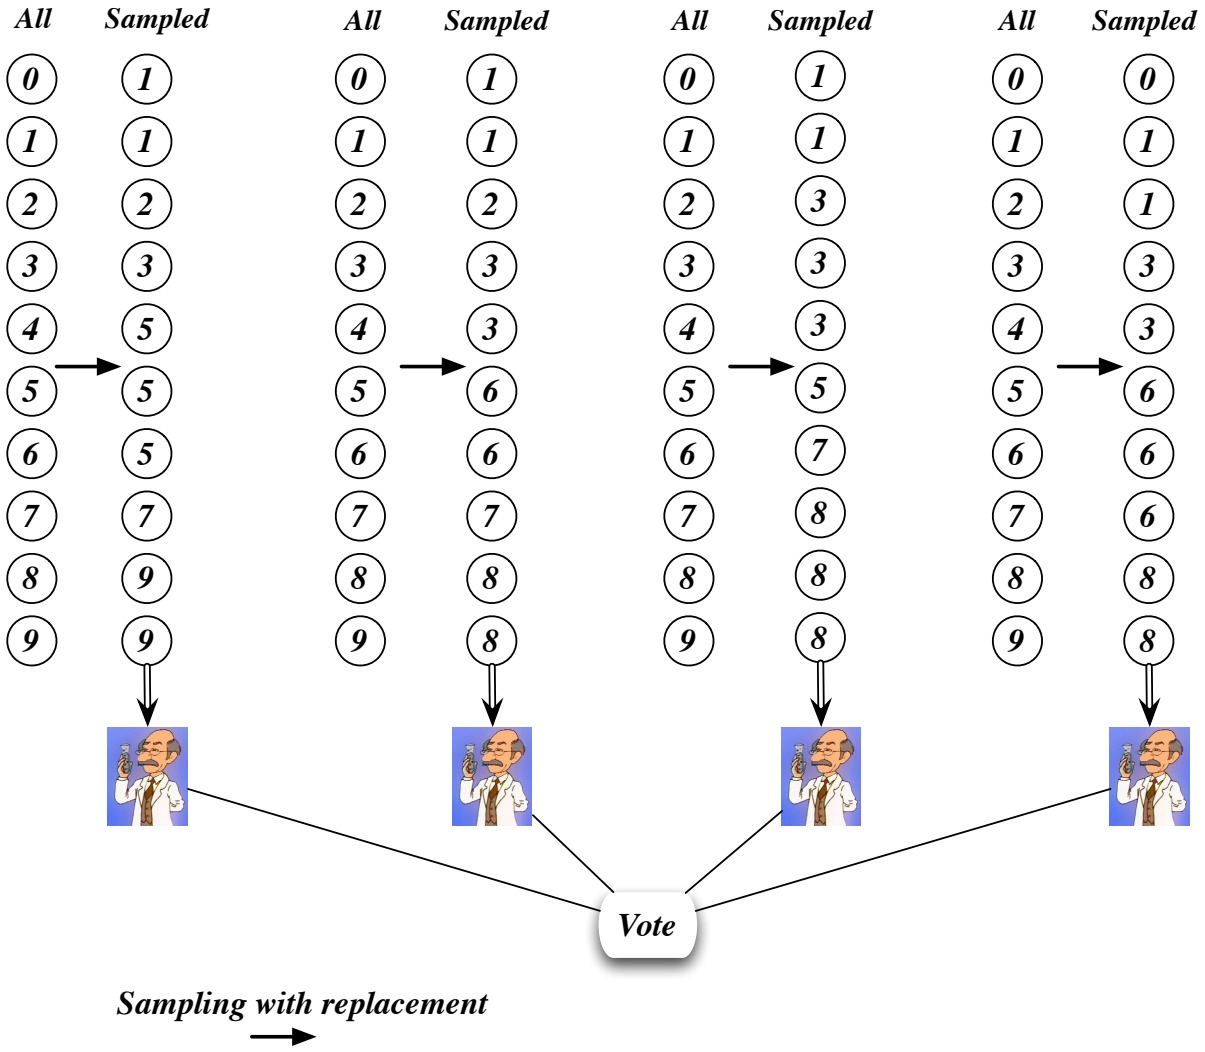

How Big An Ensemble Do You Need?

Don't use fixed size ensembles. They will short-change you and deceive you.
Instead, stop when accuracy levels off.

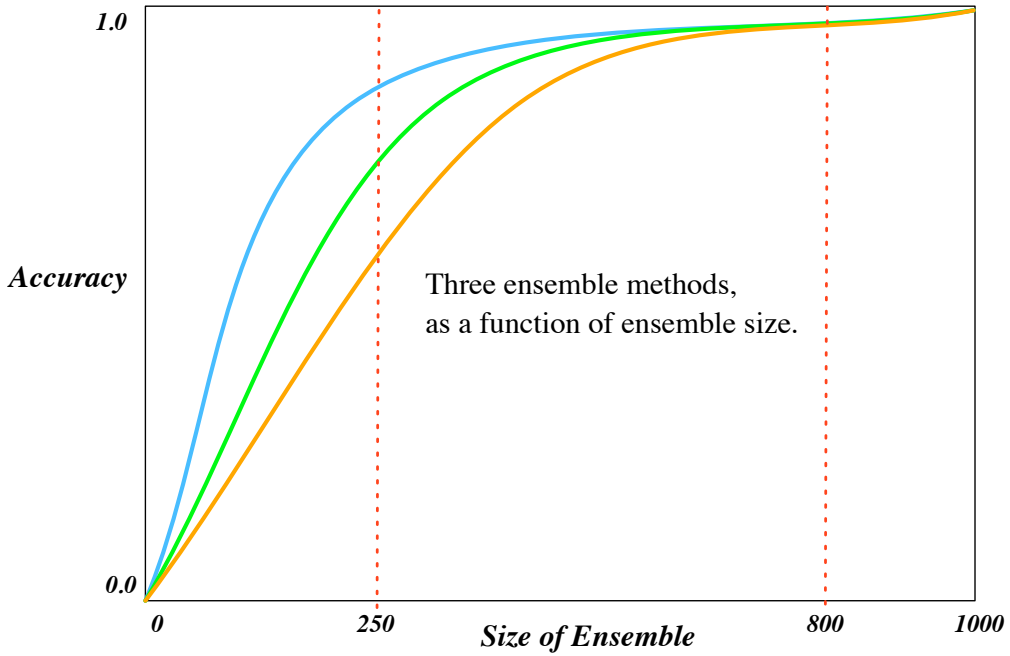

But how to measure accuracy? *Don't* just use the training data.
Use a separate validation set? Sure, but they are rare and costly.
Out-of-bag (OOB) validation is easy and cheap.

Every Classifier Lacks a Fraction of the Samples

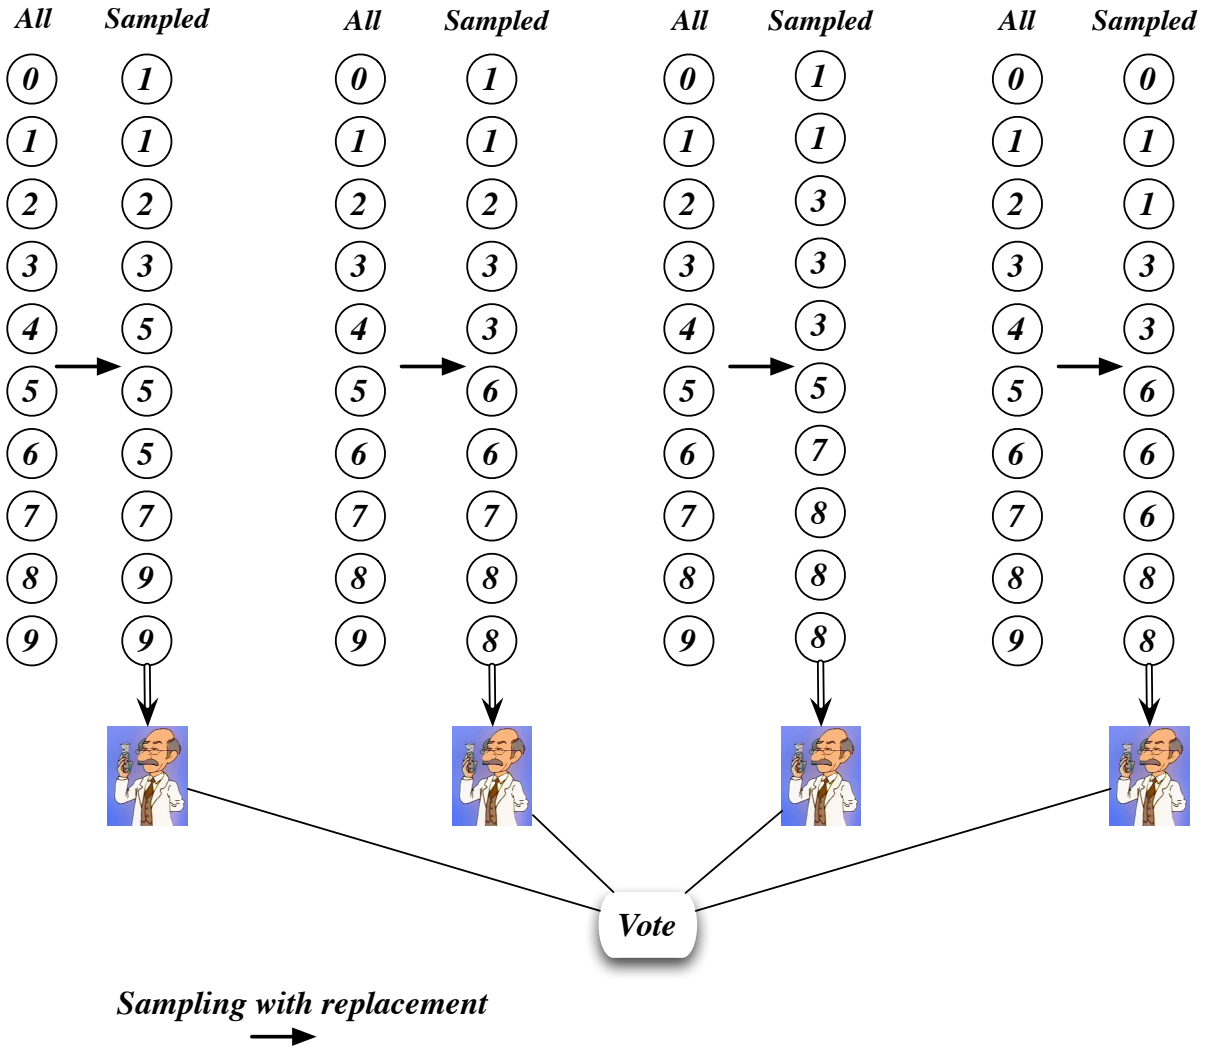

Every Sample Lacks a Fraction of the Classifiers!!

The classifiers that didn't see the sample can be fairly used to test it.

Sample 2 can be tested by E3 and E4; Sample 4 by E1, E2, E3 and E4.
Each sample can be tested by a substantial fraction of the classifiers.
So the over all accuracy is accumulated, one sample at a time.

Accuracy Increases Erratically With Ensemble Size

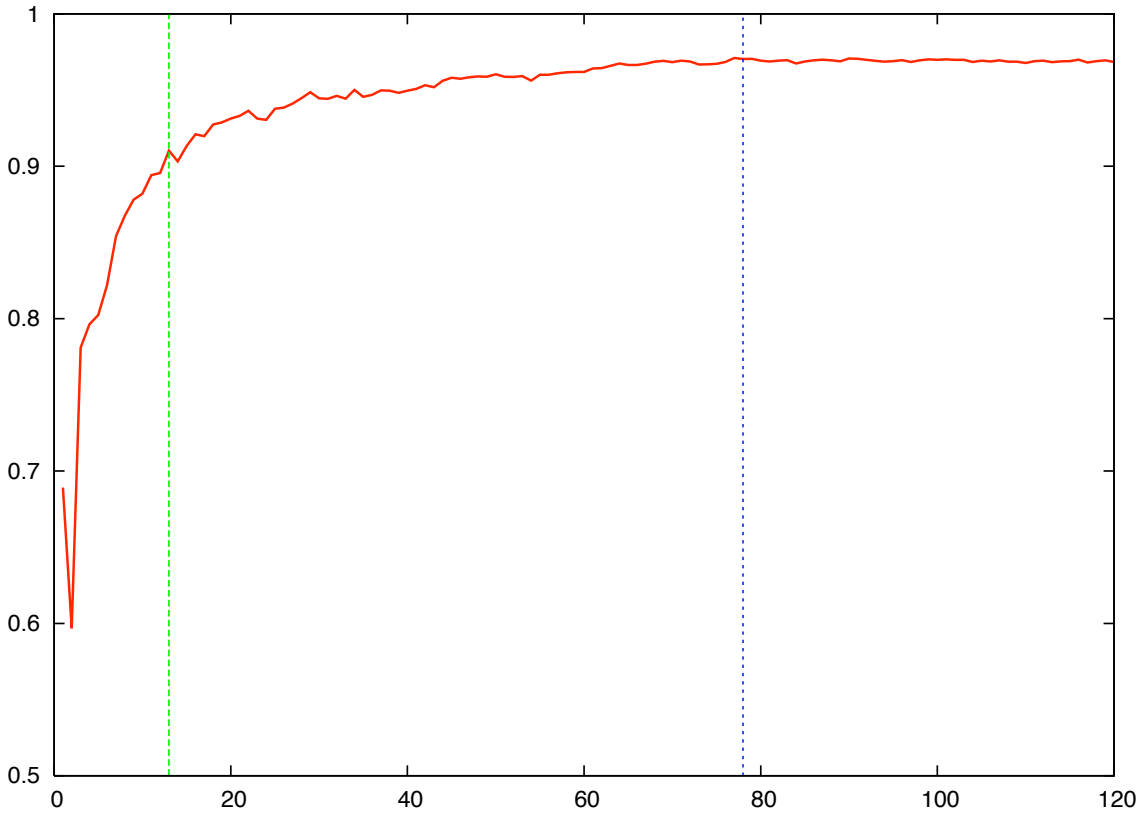

Can't stop at first peak or plateau; accuracy curve must be smoothed.

Smoothing and Selecting Via Local Windows

Three step algorithm for selecting a stop point[2]:

1. Maintain a running average over w_{small} samples, to smooth.
2. Track maximum accuracy over windows of size w_{large} until it doesn't increase.
3. Return size of ensemble that first achieved that accuracy.

Note that OOB will over-estimate ensemble size.

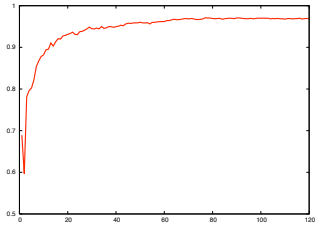

Raw Accuracy Curve

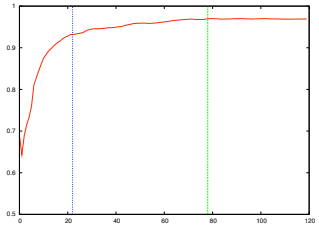

Smoothed Accuracy

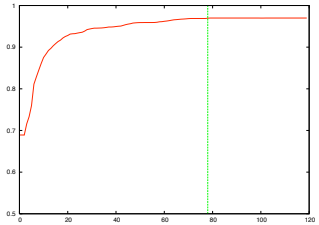

Maximum Filter Accuracy

So: Smoothed Maximum Accuracy is Effective . . .

. . . but theoretically unsatisfying.

Next Steps:

- Generate a menagerie of real curves; build intuition.
- Estimate parameters from the curve itself?
  - Extract a measure of variability from the small ensemble data?
  - Explicitly model the "noise", the variation in accuracy?
- Consult with a trained 1D signal processor.

References

[1] Banfield, R. E., Hall, L. O., Bowyer, K. W., Bhadoria, D., Kegelmeyer, W. P., and Eschrich, S. A comparison of ensemble creation techniques. In *Proceedings of the Fifth International Conference on Multiple Classifier Systems, MCS2004* (2004), J. K. F. Roli and T. Windeatt, Eds., vol. 3077 of *Lecture Notes in Computer Science*, Springer-Verlag.

[2] Banfield, R. E., Hall, L. O., Bowyer, K. W., and Kegelmeyer, W. P. A comparison of decision tree ensemble creation techniques. *IEEE Transactions on Pattern Analysis and Machine Intelligence 29*, 1 (January 2007), 173–180.