

Machine learning methods for bioinformatics

INFO-F-528

Gianluca Bontempi

Département d'Informatique
Boulevard de Triomphe - CP 212
<http://mlg.ulb.ac.be>

Bagging

- Consider a dataset D_N and a learning procedure to build an hypothesis α_N from D_N .
- The idea of **bagging** or **bootstrap aggregating** is to imitate the stochastic process underlying the realization of D_N .
- A set of B repeated bootstrap samples $D_N^{(b)}$, $b = 1, \dots, B$ are taken from D_N .
- A model $\alpha_N^{(b)}$ is built for each $D_N^{(b)}$.
- A final predictor is built by aggregating the B models $\alpha_N^{(b)}$.

R: Bagging against overfitting

- Consider a dataset $D_N = \{x_i, y_i\}$, $i = 1, \dots, N$ of $N = 100$ i.i.d. normally distributed inputs $\mathbf{x} \sim \mathcal{N}([0, 0, 0], I)$.
- Suppose that \mathbf{y} is linked to \mathbf{x} by the input/output relation

$$y = x_1^2 + 4 \log(|x_2|) + 5x_3 + \mathbf{w}$$

where $\mathbf{w} \sim \mathcal{N}(0, 0.25)$ represents the noise.

- Let's train a single-hidden-layer neural network with $s = 15$ hidden neurons on the training set.
- The prediction accuracy on the test set ($N_{\text{ts}} = 100$) is $\widehat{\text{MISE}}_{\text{ts}} = 54.43$.

- Let us apply a bagging combination with $B = 50$ (R-file `bagging.R`).
- The prediction accuracy on the test set of the bagging predictor is $\widehat{\text{MISE}}_{\text{ts}} = 6.68$. This shows that the bagging combination reduces the overfitting of the single neural network.
- Below there is the histogram of the $\widehat{\text{MISE}}_{\text{ts}}$ accuracy of each bootstrap repetition. We can see that the performance of the bagging predictor is much better than the average performance.



Boosting methods

- Boosting is one of the most powerful learning ideas introduced in the last ten years.
- Boosting is a general method which attempts to *boost* the accuracy of any given learning algorithm.
- It was originally designed for classification problems but it can profitably be extended to regression as well.
- Boosting [Freund Schapire1996,Schapire1990] encompasses a family of methods.
- The focus of boosting methods is to produce a series of “weak” learners in order to produce a powerful combination.
- A *weak learner* is a learner that has accuracy only slightly better than chance.

Boosting methods (II)

- The training set used for each member of the series is chosen based on the performance of the earlier classifier(s) in the series.
- Examples that are incorrectly predicted by previous classifiers in the series are chosen more often than examples that were correctly predicted.
- Thus Boosting attempts to produce new classifiers that are better able to predict examples for which the current ensemble's performance is poor.
- Unlike Bagging, the resampling of the training set is **dependent** on the performance of the earlier classifiers.
- The two most important types of boosting algorithms are the Ada Boost (Adaptive Boosting) algorithm (Freund, Schapire, 1997) and the Arcing algorithm (Breiman, 1996).

The Ada Boost algorithm

- Consider a classification problem where the output can take as value one of the two classes $\{-1, 1\}$. Let D_N be the training set.
- A classifier is a predictor $h(\cdot)$ which given an input x , produces a prediction taking one of the values $\{-1, 1\}$.
- A *weak classifier* is one whose misclassification error rate is only slightly better than random guessing.
- The purpose of boosting is to sequentially apply the weak classification algorithm to repeatedly modified versions of the data, thereby producing a sequence of classifiers $h_j(\cdot)$, $j = 1, \dots, m$.
- The predictions of the m *weak* classifiers are then combined through a **weighted majority vote** to produce the final prediction

$$h_{\text{boo}} = \text{sign} \left(\sum_{j=1}^m \alpha_j h_j(x, \alpha_N) \right)$$

The Ada Boost algorithm (II)

- The weights α_j of the different classifiers are computed by the algorithm. The idea is to give higher influence to the more accurate classifiers in the sequence.
- At each step, the boosting algorithm samples N times from a distribution w on the training set which put a weight w_i on each sample (x_i, y_i) , $i = 1, \dots, N$ of D_N
- Initially the weights are all set to $w_i = 1/N$ so that the first step simply trains the classifier in the standard manner.
- For each successive iteration $j = 1, \dots, m$ the probability weights are individually modified and the classification algorithm is re-applied to the resampled training set.

The Ada Boost algorithm (III)

- At a generic step j these observations that were misclassified by the classifier $h_{j-1}(\cdot)$ trained at the previous step, have their weights w_i increased, whereas the weights are decreased for those that were classified correctly.
- The rationale is that, as iterations proceed, observations that are difficult to correctly classify receive ever-increasing influence and the classifier is forced to concentrate on them.
- Note the presence in the algorithm of **two types of weights**: the weights $\alpha_j, j = 1, \dots, m$ that measure the importance of the classifiers and the weights $w_i, i = 1, \dots, N$ of the samples.
- Weak learners are added until some desired low training error has been achieved.

Ada Boost in detail

- 1 Initialize the observation weights $w_i = 1/N$, $i = 1, \dots, N$.
- 2 For $j = 1$ to m :
 - 1 Fit a classifier $h_j(\cdot)$ to the training data obtained by resampling D_N using weights w_i .
 - 2 Compute the misclassification error on the training set

$$\widehat{\text{MME}}_{\text{emp}}^{(j)} = \frac{\sum_{i=1}^N w_i I(y_i \neq h_j(x_i))}{\sum_{i=1}^N w_i}$$

- 3 Compute

$$\alpha_j = \log\left(\frac{1 - \widehat{\text{MME}}_{\text{emp}}^{(j)}}{\widehat{\text{MME}}_{\text{emp}}^{(j)}}\right)$$

Note that $\alpha_j > 0$ if $\widehat{\text{MME}}_{\text{emp}}^{(j)} \leq 1/2$ (which we can assume without loss of generality) and that α_j gets larger as $\widehat{\text{MME}}_{\text{emp}}^{(j)}$ gets smaller.

Ada Boost in detail (II)

- 2 4 For $i = 1, \dots, N$ set

$$w_i \leftarrow w_i \begin{cases} \exp[-\alpha_j] & \text{if correctly classified} \\ \exp[\alpha_j] & \text{if incorrectly classified} \end{cases}$$

- 5 The weights are normalized to ensure that w_i represents a true distribution.
- 3 Output of the weighted majority vote

$$h_{\text{boo}} = \text{sign} \left(\sum_{j=1}^m \alpha_j h_j(x, \alpha_N) \right)$$

R TP: Ada Boost for classification

- Consider the medical dataset *Pima* obtained by a statistical survey on women of Pima Indian heritage. This dataset reports the presence of diabetes in Pima indian women together with other clinical measures (blood pressure, insulin, age,...).
- The classification task is to predict the presence of diabetes as a function of the clinical measures.
- We consider a training set of $N = 40$ and a test set of 160 samples.
- The classifier is a simple classification tree which returns a misclassification rate $\widehat{MME}_{ts} = 0.36$.
- We use a boosting procedure with $m = 15$ to improve the performance of the weak classifier. The misclassification rate of the boosted classifier is $\widehat{MME}_{ts} = 0.24$.
- R file `boosting.r`.

Theoretical foundation of boosting

- Boosting has its roots in a theoretical framework for studying machine learning called the PAC learning model.
- Freund and Scapire proved that the empirical error of the final hypothesis h_{boo} is at most

$$\prod_{j=1}^m \left[2\sqrt{\widehat{\text{MME}}_{\text{emp}}^{(j)} * (1 - \widehat{\text{MME}}_{\text{emp}}^{(j)})} \right]$$

- They showed also how to bound the generalization error.

The arcing algorithm

- This was proposed as a modification of the original Ada Boost algorithms by Breiman.
- It is based on the idea that the success of boosting is related to the adaptive resampling property where increasing weight is placed on those samples more frequently misclassified.
- ARCing stays for *Adaptive Resampling and Combining*.
- The complex updating equations of Ada Boost are replaced by much simpler formulations.
- The final classifier is obtained by unweighted voting

ARcing in detail

- 1 Initialize the observation weights $w_i = 1/N$, $i = 1, \dots, N$.
- 2 For $j = 1$ to m :
 - 1 Fit a classifier h_j to the training data obtained by resampling D_N using weights w_i .
 - 2 Let e_j be the number of misclassifications of the i th sample by the j classifiers h_1, \dots, h_j .
 - 3 The updated weights are defined by

$$w_i = \frac{1 + e_i^4}{\sum_{i=1}^N (1 + e_i^4)}$$

- 3 The output is obtained by unweighted voting of the m classifiers h_j .

R example: ARcing for classification

- Consider the medical dataset *Breast Cancer* obtained by Dr. William H. Wolberg (physician) at the University of Wisconsin Hospital in USA. This dataset reports the class of cancer (malignant and benign) together with other properties (clump thickness, uniformity of cell size, uniformity of cell shape, marginal adhesion,...)
- The classification task is to predict the class of the breast cancer as a function of the clinical measures.
- We consider a training set of $N = 400$ and a test set of 299 samples.
- The classifier is a simple classification tree which returns a misclassification rate $\widehat{\text{MME}}_{\text{ts}} = 0.063$.
- We use an arcing procedure with $m = 15$. It gives a misclassification rate $\widehat{\text{MME}}_{\text{ts}} = 0.016$.
- R file `arcing.r`

Some considerations

- Boosting is a very recent technique. According to most researchers, it is one of the most promising techniques. However, it is probably too early to have a thorough assessment of the technique.
- Boosting techniques are very simple and easy to program.
- They have few parameters to tune.
- They advocate a shift in the attitude of the learning-system designer: instead of trying to design a learning algorithm which should be accurate over the entire space, he can instead focus on finding weak algorithms that only need to be better than random.
- A nice property of Ada Boost is its ability to identify outliers (hard samples).
- As for bagging, the main effect of boosting is to reduce variance. It works effectively for high variance classifiers.
- Boosting seems to do better than bagging (see Breiman's work)

Some caveats

- The actual performance of boosting on a particular problem is clearly dependent on the data and the weak learner.
- Boosting can fail to perform well given
 - insufficient data,
 - overly complex weak hypothesis
 - weak hypothesis that are too weak.