

The Pennsylvania State University
The Graduate School

A MULTICLASS BOOSTING CLASSIFICATION METHOD
WITH ACTIVE LEARNING

A Thesis in
Information Sciences and Technology
by
Jian Huang

© 2009 Jian Huang

Submitted in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science

December 2009

The thesis of Jian Huang was reviewed and approved* by the following:

C. Lee Giles
Professor of Information Sciences and Technology
Thesis Advisor

Prasenjit Mitra
Assistant Professor of Information Sciences and Technology

Jia Li
Associate Professor of Statistics

Henry C. Foley
Professor of Information Sciences and Technology
Dean of the College of Information Sciences and Technology

*Signatures are on file in the Graduate School.

Abstract

Boosting methods, a class of ensemble learning methods, have been very popular in the machine learning and data mining fields. The classical AdaBoost algorithm only requires the underlying weak learner to perform better than random guessing and has shown empirical success in its resistance to overfitting. Many challenges, however, remain to be tackled. There are a number of theoretical research questions such as how to efficiently handle the multiclass classification setting, as well as how to reduce the impacts of outliers to the ensemble model. From an empirical perspective, scaling up a multiclass boosting algorithm, which requires the training of many underlying weak learners, to a large scale dataset is of practical interest.

A novel multiclass classification algorithm Gentle Adaptive Multiclass Boosting Learning (GAMBLE) is proposed to address these issues. The algorithm naturally extends the two class Gentle AdaBoost algorithm to multiclass classification by using the multiclass exponential loss and the multiclass response encoding scheme. Unlike other multiclass algorithms which reduce the K -class classification task to K binary classifications, GAMBLE handles the task directly and symmetrically, with only one committee classifier. We formally derive the GAMBLE algorithm with the quasi-Newton method, and prove the structural equivalence of the two regression trees in each boosting step.

To scale up to large datasets, we utilize the generalized Query By Committee (QBC) active learning framework to focus learning on the most informative samples. Empirical results show that with QBC-style active sample selection, faster training time and potentially higher classification accuracy can be achieved by using only the informative fraction of the training instances. GAMBLE's numerical superiority, structural elegance and low computation complexity make it highly competitive with state-of-the-art multiclass classification algorithms.

Table of Contents

List of Figures	vi
List of Tables	vii
List of Algorithms	viii
Acknowledgments	ix
Chapter 1	
Introduction	1
Chapter 2	
Related Work	3
2.1 Related research in multiclass boosting	3
2.2 Related work in active learning	4
Chapter 3	
The GAMBLE Boosting Algorithm	5
3.1 The General GAMBLE Algorithm	5
3.2 The GAMBLE Algorithm Using Weighted Regression Trees	8
3.2.1 Weighted Regression Trees	8
3.2.2 GAMBLE with Weighted Regression Trees	9
3.3 Remarks on the GAMBLE Algorithm	11
Chapter 4	
Query By Committee Style Active Learning	13
Chapter 5	
Experiments	15
5.1 Experiments on Real World Data	16
5.2 Active Learning Performance	18
5.3 Simulation Study	21

Chapter 6	
Conclusions	24
Appendix A	
Proofs of Theorems	25
A.1 Proof of Theorem 2	25
A.2 Proof of Corollary 1	27
Appendix B	
Proofs of Results	28
B.1 Proof of Result 2	28
B.2 Proof of Result 3	28
Bibliography	29

List of Figures

3.1	The concave transformation of the regression function $g(x)$ to the (un-centralized) weak learner $r(x)$	10
5.1	GAMBLE boosting regression trees of different sizes in the Pendigits dataset. T denotes the number of leaves in the regression tree. Note that the iteration number is shown on a log scale.	16
5.2	Test error scatter plot for GAMBLE and SAMME at 10, 100 and 1,000 iterations in 8 UCI multiclass datasets. Each point shows the error rate of GAMBLE and SAMME on one benchmark dataset.	17
5.3	Test error (in the range of $[0, 1]$) for GAMBLE and SAMME.Discrete in the <i>Glass</i> dataset.	17
5.4	Test error (in the range of $[0, 1]$) for GAMBLE and SAMME.Discrete in the <i>Segmentation</i> dataset.	18
5.5	Test error of Active GAMBLE and Random GAMBLE in the <i>Pendigits</i> dataset. Active GAMBLE incurs lower test error and uses fewer training samples than Random GAMBLE. Note that the test error curves for the first 1000 samples are zoomed in for the <i>Pendigits</i> dataset.	19
5.6	Test error of Active GAMBLE and Random GAMBLE in the <i>Satimage</i> dataset. Active GAMBLE incurs lower test error and uses fewer training samples than Random GAMBLE.	19
5.7	GAMBLE's training and test errors in the concentric spheres example ($K=2$). Iteration number is shown on the log scale, highlighting the significant difference in training error iterations to those of testing.	22
5.8	GAMBLE's training and test errors in the concentric spheres example ($K=5$). Iteration number is shown on the log scale, highlighting the significant difference in training error iterations to those of testing.	22

List of Tables

- 5.1 UCI benchmark datasets. K is the number of classes and D is the number of feature dimensions. The last two columns are the numbers of the training and testing instances, respectively. 15
- 5.2 Data efficiency comparison. *Lowest error rate* is the minimum error rate incurred by Active GAMBLE. *Target error rate* is the lowest error rate incurred by Batch GAMBLE in the first 1,000 iterations. 20
- 5.3 Training time comparison. 20

List of Algorithms

1	The Gentle Adaptive Multiclass Boosting Learning (GAMBLE) Algorithm. . . .	7
2	Gentle Adaptive Multiclass Boosting Learning (GAMBLE) Using Weighted Regression Trees.	10
3	Gentle AdaBoost (Friedman et al. [8]).	11
4	Active GAMBLE Algorithm.	13

Acknowledgments

I am deeply indebted to my advisor Prof. C. Lee Giles, who showed me the beauty of science and guided and supported my academic pursuits. I would like to thank Prof. Hongyuan Zha for the in-depth feedback and suggestions for this work. I would also like to thank Dr. Seyda Ertekin for her input and our fruitful discussions on the active learning topic. I am grateful for the helpful comments from my committee members, Prof. Jia Li and Prof. Prasenjit Mitra, as well as the anonymous reviewers of this work.

Dedication

This thesis is specially dedicated to my parents for their encouragement for my academic pursuits and their invaluable support in life.

Chapter 1

Introduction

Boosting algorithms [16, 13] have been very popular in the past decade. The general boosting framework only requires the underlying weak learners to perform better than random guessing. By combining these weak learners based on weighted majority vote, a committee classifier dramatically reduces the training and testing error rates. The representative AdaBoost algorithm [5, 8], in particular, adapts to the weighted error of the current weak learner, by strengthening the weights of the misclassified samples and dampening the weights of the correctly classified ones.

The most appealing feature of boosting is its empirical *resistance to overfitting* for various classification tasks. Researchers have gained insights into this seemingly mysterious phenomenon. Friedman et al. [8] provided an alternative viewpoint for boosting as the well-known forward stagewise additive modeling. By combining low-variance (high-bias) weak learners, boosting acts as a bias reduction process. Based on the margin and VC dimension theory, Schapire et al. [5] proved that the generalization error is upper bounded, independent of the number of training iterations. Boosting is thus regarded as a margin increasing process, which connects it to Support Vector Machines (SVM) [20, 21] which explicitly maximizes the margin.

This work is motivated by two goals. First, boosting methods were initially designed to handle binary classification problems. Initial efforts typically decomposed the K -class classification problem to K binary classification problems by using the ‘one-against-all’ strategy, thus requiring K committees to be trained. A better approach should handle such problems directly and symmetrically. Since weak learners (e.g. classification trees) can naturally handle multiclass cases, a direct multiclass boosting method can unleash the innate power of multiclass weak learners that a ‘one-against-all’ or ‘pairwise’ method cannot. Generalizing to the multiclass cases, however, requires non-trivial extension of the loss function, the encoding scheme and the algorithms per se.

Furthermore, multiclass labels are usually expensive to obtain (e.g. text classification). And even when such training data is abundant, multiclass classification is generally less efficient.

The practicality of batch learning algorithms is limited in such learning settings. Boosting methods, specifically, require global weight updates in each boosting step and multiple iterations to converge, as such they are not particularly efficient for large-scale multiclass learning. Active learning strategies enable large-scale learning for boosting algorithms by greatly improving data efficiency.

This work proposes a novel multiclass boosting algorithm, Gentle Adaptive Multiclass Boosting Learning (GAMBLE) [11], which addresses the aforementioned issues and has the following features:

- A K -class classification problem is treated simultaneously and symmetrically, without recasting it to multiple binary classification problems.
- Only one regression function per iteration is fitted, yielding a highly efficient solution.
- The algorithm is less sensitive to outliers in training and robust to unseen data in testing, due to its numerical stability in each weak learner.
- With active sample selection, Active GAMBLE uses the informative portion of the training samples and is thus fast in training with early stopping. It is also efficient in prediction with a simpler model and potentially has greater generalization power.
- The algorithmic structure is succinct and the algorithm is easy to implement.

The thesis is organized as follows. The related research on multiclass boosting algorithms and active learning is reviewed in Chapter 2. In Chapter 3, we first derive the general GAMBLE algorithm using quasi-Newton methods and then propose an efficient implementation of GAMBLE using regression trees. We also comment on the advantages of GAMBLE over other multiclass boosting algorithms. In Chapter 4, we further investigate the active learning strategy for GAMBLE, which enables large-scale learning, minimizes labeling and improves both accuracy and time complexity. We present experimental results on real world and synthetic data in Chapter 5 and conclude in Chapter 6.

Related Work

2.1 Related research in multiclass boosting

The first attempt to extend AdaBoost to multiclass classification is the AdaBoost.M1 algorithm [5], a direct generalization of AdaBoost by replacing the error with an indicator function $I[h_k(\mathbf{x}_i) \neq y_i]$, where h_k is the k -th dimension of the weak hypothesis. Since the same binary loss function is used as AdaBoost.Discrete in the binary case, AdaBoost.M1 fails when the underlying weak learners misclassify more than half of the samples. This is not appropriate for multiclass classification, when the underlying weak learners perform slightly better than random guessing ($err < \frac{K-1}{K}$).

More sophisticated approaches have been investigated for this case. AdaBoost.M2 [5, 15] (a special case of AdaBoost.MR [18]) solves the multiclass problem by using the ‘*pairwise*’ strategy. However, AdaBoost.M2 complicates the design of the weak learners since it requires the optimization of weak hypotheses for the pseudoloss measure, a weighted transform of the weak hypotheses. AdaBoost.MH [18] was proposed as an alternative to AdaBoost.M2. This method adopts the well-known ‘*one-against-all*’ [7] strategy to recast the K -class classification problem to K binary classification problems, which are in turn solved by the binary AdaBoost.Real [8] algorithm. Other multiclass variants exist such as AdaBoost.MO [18, 2] which uses the correcting output code.

These boosting algorithms do not handle the multiclass classification problems directly, but instead reduce them to multiple two-class classification problems. In practice, even when the classification boundaries for class pairs are simple, the pooling classes (as in the ‘against all’ classes) may form boundaries that are too complex to approximate and the methods above may fail [8, 7, 9]. Friedman et al. [8] proposed the multiclass LogitBoost algorithm by using the multiclass logistic model, which outperformed the asymmetric variants such as AdaBoost.MH. This gave insight into how to solve multiclass classification problems simultaneously. Zhu et. al. [22] proposed the multiclass exponential loss function. By adopting the K -class response encoding

scheme, they successfully generalized AdaBoost.Discrete and AdaBoost.Real to the multiclass version SAMME.Discrete and SAMME.Real. SAMME (Stagewise Additive Modeling using a Multi-class Exponential loss function) generates only one K -class classifier in each iteration, thus it is K times faster than AdaBoost.MH.

As noted by Schapire et. al. [16] and convincingly shown by Dietterich [4], when there are a great number of outliers in the training instances, AdaBoost (as well as SAMME.Discrete and SAMME.Real) may overemphasize them, which eventually leads to inferior hypotheses. Friedman et. al. proposed the Gentle AdaBoost algorithm [8] to ameliorate such a problem, which is a gradient ascent approach in function space. Gentle AdaBoost also outperforms LogitBoost [8] for its numerical stability, since the updates are bounded. The Gentle AdaBoost.MH algorithm was also proposed in [8] but still used the ‘*one-against-all*’ strategy. This work generalizes this successful boosting algorithm to the multiclass case in non-trivial ways, yielding a natural and efficient solution.

2.2 Related work in active learning

Compared to the profusion of literature in boosting algorithms, there is relatively little work on the scalability issue of boosting methods. Active learning strategies can be used to scale up boosting algorithms by focusing training on the most informative samples. Query By Committee (QBC) [19] is a popular and theoretically well motivated active learning method. Abe et al. [1] proposed the QBoost algorithm which combines QBC and AdaBoost to efficiently learn the committee. This work adopts the same active sample selection principle as in QBoost and extends it to handle multiclass classification.

The GAMBLE Boosting Algorithm

We first propose the general GAMBLE algorithm with theoretical justification. Based on the weighted least squares criterion and the regression tree method, we then propose an efficient implementation of the GAMBLE algorithm which only fits one regression function to the true response. Additionally, we discuss the numerical and computational advantages of GAMBLE over other multiclass boosting alternatives.

Before proceeding to the methods, we first formalize the K -class boosting learning problem as follows:

Definition 1. Given N training samples $\{(\mathbf{x}_i, c_i)\}_{i=1}^N$ ($\mathbf{x}_i \in \mathcal{R}^D, c_i \in \{1, 2, \dots, K\}$) drawn independently from the population, we need to obtain M weak hypotheses $\{\mathbf{f}^{(m)}(\mathbf{x}) | \mathbf{f}^{(m)} \in \mathcal{H}\}_{m=1}^M$, each performing better than random guessing ($\exists \gamma > 0, \text{err}(\mathbf{f}^{(m)}) \leq \frac{K-1}{K} - \gamma$). The combined committee $\mathbf{F}(\mathbf{x})$ constitutes a strong hypothesis, in terms of minimizing the empirical loss.

In what follows, we use the K -class response encoding scheme as in [12, 22] rather than a discrete number as a class label. A class label c is encoded by a K -dimensional response vector $\mathbf{y} = (y_1, \dots, y_K)^T$, where $y_k = 1$ if and only if $k = c$ and $y_k = -\frac{1}{K-1}$ otherwise. Hence, the response vector satisfies the sum-to-zero constraint, i.e. $\sum_{k=1}^K y_k = 0$. The advantage of this encoding scheme becomes more clear in the sequel.

3.1 The General GAMBLE Algorithm

Given a classifier \mathbf{F} and a sample (\mathbf{x}, \mathbf{y}) , a loss function $J(\mathbf{F}(\mathbf{x}))$ measures the degree of disappointment in the difference between the true response and the learner's prediction. To generalize to the multiclass case, we adopt the multiclass AdaBoost exponential loss [22], $J(\mathbf{F}(\mathbf{x})) = \exp(-\frac{1}{K}\mathbf{y}^T\mathbf{F}(\mathbf{x}))$, a strictly convex function that has successfully generalized the original AdaBoost algorithm to the multiclass case. Our goal is to minimize the empirical loss. To this end, we use quasi-Newton steps to derive the multiclass extension of the Gentle AdaBoost algorithm, as shown in the following result.

Result 1. *The GAMBLE algorithm, in the population version, uses quasi-Newton steps for minimizing the multiclass exponential loss $E[\exp(-\frac{1}{K}\mathbf{y}^T\mathbf{F}(\mathbf{x}))|\mathbf{x}]$.*

Proof. The goal here is to derive the quasi-Newton update $\mathbf{f} = (f_1, \dots, f_K)$.

Given the current committee $\mathbf{F}(\mathbf{x})$, we calculate the improved update $\mathbf{f}(\mathbf{x})$. Consider the expected multiclass exponential loss J ,

$$\begin{aligned} & E[J(\mathbf{F}(\mathbf{x}) + \mathbf{f}(\mathbf{x}))|\mathbf{x}] \\ &= E[\exp(-\frac{1}{K}\mathbf{y}^T(\mathbf{F}(\mathbf{x}) + \mathbf{f}(\mathbf{x})))|\mathbf{x}] \end{aligned} \quad (3.1)$$

- (i) Conditioning on input \mathbf{x} , we compute the first and the second order derivatives of the exponential loss at $\mathbf{f}(\mathbf{x}) = \mathbf{0} = (0, 0, \dots, 0)^T$. The first order derivative with respect to the j -th dimension is:

$$\begin{aligned} s_j(\mathbf{x}) &= \left. \frac{\partial(J(\mathbf{F}(\mathbf{x}) + \mathbf{f}(\mathbf{x})))}{\partial f_j(\mathbf{x})} \right|_{\mathbf{f}(\mathbf{x})=\mathbf{0}} \\ &= E\left[-\frac{1}{K}y_j \exp(-\frac{1}{K}\mathbf{y}^T(\mathbf{F}(\mathbf{x}) + \mathbf{f}(\mathbf{x})))|\mathbf{x}\right] \Big|_{\mathbf{f}(\mathbf{x})=\mathbf{0}} \\ &= E\left[-\frac{1}{K}y_j \exp(-\frac{1}{K}\mathbf{y}^T\mathbf{F}(\mathbf{x}))|\mathbf{x}\right] \end{aligned}$$

Likewise, the Hessian is:

$$\begin{aligned} H_{j,k}(\mathbf{x}) &= \left. \frac{\partial^2(J(\mathbf{F}(\mathbf{x}) + \mathbf{f}(\mathbf{x})))}{\partial f_j(\mathbf{x})\partial f_k(\mathbf{x})} \right|_{\mathbf{f}(\mathbf{x})=\mathbf{0}} \\ &= E\left[\frac{1}{K^2}y_j y_k \exp(-\frac{1}{K}\mathbf{y}^T\mathbf{F}(\mathbf{x}))|\mathbf{x}\right] \end{aligned}$$

- (ii) The quasi-Newton update is used to reduce the substantial computation overhead of Newton's method, the former of which uses the diagonal approximation of the Hessian, i.e.

$$f_j(\mathbf{x}) \leftarrow K \cdot \frac{E[y_j \exp(-\frac{1}{K}\mathbf{y}^T\mathbf{F}(\mathbf{x}))|\mathbf{x}]}{E[y_j^2 \exp(-\frac{1}{K}\mathbf{y}^T\mathbf{F}(\mathbf{x}))|\mathbf{x}]} \quad (3.2)$$

- (iii) Using the **weighted conditional expectation** notation [8],

$$E_w[\mathbf{g}(\mathbf{x}, \mathbf{y})] \stackrel{\text{def}}{=} \frac{E[w(\mathbf{x}, \mathbf{y})\mathbf{g}(\mathbf{x}, \mathbf{y})|\mathbf{x}]}{E[w(\mathbf{x}, \mathbf{y})|\mathbf{x}]} \quad (3.3)$$

where weight $w = w(\mathbf{x}, \mathbf{y}) = \exp(-\frac{1}{K}\mathbf{y}^T\mathbf{F}(\mathbf{x}))$,

(3.2) can be rewritten as

$$\begin{aligned} f_j(\mathbf{x}) &\leftarrow K \cdot \frac{E[y_j w(\mathbf{x}, \mathbf{y})|\mathbf{x}]}{E[y_j^2 w(\mathbf{x}, \mathbf{y})|\mathbf{x}]} \\ &= K \cdot \frac{E[y_j w(\mathbf{x}, \mathbf{y})|\mathbf{x}]/E[w(\mathbf{x}, \mathbf{y})|\mathbf{x}]}{E[y_j^2 w(\mathbf{x}, \mathbf{y})|\mathbf{x}]/E[w(\mathbf{x}, \mathbf{y})|\mathbf{x}]} \end{aligned}$$

Therefore, we express the update as

$$f_j(\mathbf{x}) \leftarrow K \cdot \frac{E_w[y_j|\mathbf{x}]}{E_w[y_j^2|\mathbf{x}]}. \quad (3.4)$$

□

Algorithm 1 The Gentle Adaptive Multiclass Boosting Learning (GAMBLE) Algorithm.

Input: Training data set $S = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$

- 1: Initialization: set the observation weights $w_i \leftarrow \frac{1}{N}$, $i = 1, 2, \dots, N$, and set $\mathbf{F}(\mathbf{x}) \leftarrow 0$.
- 2: **for** $m \leftarrow 1$ to M **do**
- 3: Using weight distribution $\{w_i\}$, fit two regression functions $\mathbf{g}^{(m)}(\mathbf{x})$ of \mathbf{y} and $\mathbf{h}^{(m)}(\mathbf{x})$ of \mathbf{z} to \mathbf{x} by weighted least-squares (WLS), where $z_j = y_j^2$, $j = 1, 2, \dots, K$.
- 4: Set $r_j^{(m)}(\mathbf{x}) = K \cdot \frac{g_j^{(m)}(\mathbf{x})}{h_j^{(m)}(\mathbf{x})}$
- 5: Obtain weak learner $f_j^{(m)}(\mathbf{x}) = r_j^{(m)}(\mathbf{x}) - \frac{1}{K} \sum_{k=1}^K r_k^{(m)}(\mathbf{x})$, $j = 1, 2, \dots, K$.
- 6: Update committee $\mathbf{F}(\mathbf{x}) \leftarrow \mathbf{F}(\mathbf{x}) + \mathbf{f}^{(m)}(\mathbf{x})$.
- 7: Update weights $w_i \leftarrow w_i \exp(-\frac{1}{K} \mathbf{y}_i^T \mathbf{f}^{(m)}(\mathbf{x}_i))$, $i = 1, 2, \dots, N$ and renormalize.
- 8: **end for**

Output: the committee classifier: $C(\mathbf{x}) = \arg \max_k F_k(\mathbf{x})$.

The Gentle Adaptive Multiclass Boosting Learning (GAMBLE) [11] is directly translated from Result 1 (see Algorithm 1). Note that in step (5), the weak learner is ‘centralized’ in each dimension such that it satisfies the symmetric constraint $\sum_{k=1}^K f_k^{(m)} = 0$. This does not affect the weight update due to the encoding scheme of the response, because,

$$\begin{aligned} \mathbf{y}^T \mathbf{f}^{(m)}(\mathbf{x}) &= \sum_{k=1}^K y_k f_k^{(m)} = \sum_{k=1}^K y_k r_k^{(m)} \\ &\quad - \frac{1}{K} \sum_{k=1}^K y_k \sum_{k=1}^K r_k^{(m)}(x) = \mathbf{y}^T \mathbf{r}^{(m)}(\mathbf{x}). \end{aligned}$$

Also note that our choice of multiclass exponential loss is valid, because in fact the population minimizer for the expected loss agrees with the Bayes optimal classification rule, as shown by the following theorem:

Theorem 1. With exponential loss $\exp(-\frac{1}{K} \mathbf{y}^T \mathbf{f}(x))$,

$$\arg \max_k r_k(\mathbf{x}) = \arg \max_k f_k(\mathbf{x}) = \arg \max_k Pr(c = k|\mathbf{x}).$$

Proof. The constrained minimization problem for the empirical loss can be formulated as:

$$\arg \min_{\mathbf{f}(\mathbf{x})} \mathbf{E}_{\mathbf{Y}|\mathbf{x}} \exp(-\frac{1}{K} \mathbf{Y}^T \mathbf{f}(\mathbf{x}))$$

subject to $f_1 + \dots + f_K = 0$.

The population minimizer can be found by the Lagrange multiplier (see [22] for similar derivation),

$$f_k = (K - 1) \left(\log Pr(c = k|\mathbf{x}) - \frac{1}{K} \sum_{j=1}^K \log Pr(c = j|\mathbf{x}) \right)$$

Thus $\arg \max_k f_k(\mathbf{x}) = \arg \max_k Pr(c = k|\mathbf{x})$. Obviously, $\arg \max_k r_k(\mathbf{x}) = \arg \max_k f_k(\mathbf{x})$. \square

When implemented on data, the weighted expectation in (Equation 3.4) can be substituted by weighted regression methods. As an ensemble learning algorithm, the general GAMBLE algorithm may employ an arbitrary regressor as a weak learner. Since most boosting methods work well with ‘small’ classification trees as weak learners, we implement the general GAMBLE algorithm with weighted regression trees [3] in our case.

3.2 The GAMBLE Algorithm Using Weighted Regression Trees

Boosting algorithms have been predominantly implemented using classification and regression trees. We first extend the regression tree method with the weighted least squares criterion. Then we explore the relationship of the two regression functions in the general GAMBLE algorithm, and propose an efficient implementation that requires only one regression.

3.2.1 Weighted Regression Trees

Since weak learners work with weighted samples in boosting methods, we extend the classical regression tree method [10, 3] to the weighted version with the weighted least squares (WLS) criterion.

Suppose our model partitions the entire space (spanned by the variables X_1, X_2, \dots, X_D) into J regions R_1, R_2, \dots, R_J . In each region R_j , the model produces a constant response \mathbf{c}_{R_j} as the fitted value:

$$\mathbf{g}(\mathbf{x}) = \sum_{j=1}^J \mathbf{c}_{R_j} I[\mathbf{x} \in R_j] \quad (3.5)$$

\mathbf{c}_{R_j} is chosen by minimizing the WLS criterion, i.e.,

$$\mathbf{c}_{R_j} = \arg \min_{\mathbf{c}} \sum_{\mathbf{x}_l \in R_j} w_l (\mathbf{y}_l - \mathbf{c})^T (\mathbf{y}_l - \mathbf{c}) \quad (3.6)$$

By taking the first order derivative and setting it to zero, \mathbf{c}_{R_j} is the weighted mean of the true responses,

$$c_{k,R_j} = \frac{\sum_{\mathbf{x}_l \in R_j} w_l y_{k,l}}{\sum_{\mathbf{x}_l \in R_j} w_l} \quad (k = 1, 2, \dots, K) \quad (3.7)$$

In terms of the WLS criterion, seeking an optimal binary regression tree is inhibitedly expensive. CART [3], however, grows the tree in a top-down and greedy fashion and thus dramatically reduces the search space. Specifically, only one variable is taken into account in each splitting step, thus the splitting hyperplane is always orthogonal to the axis of the corresponding variable. More precisely, each splitting step is a decision making process: the region R is split into two subregions R_1 and R_2 , where the points satisfying the decision rule $X_j \leq h$ should be grouped into R_1 and otherwise R_2 . The parameters (j, h) in the decision rule are chosen such that they also minimize the WLS criterion:

$$(j, h) = \arg \min_{(j, h)} \left[\sum_{\mathbf{x}_l \in R_1(j, h)} w_l (\mathbf{y}_l - \mathbf{c}_{R_1(j, h)})^T (\mathbf{y}_l - \mathbf{c}_{R_1(j, h)}) \right. \\ \left. + \sum_{\mathbf{x}_l \in R_2(j, h)} w_l (\mathbf{y}_l - \mathbf{c}_{R_2(j, h)})^T (\mathbf{y}_l - \mathbf{c}_{R_2(j, h)}) \right] \quad (3.8)$$

This decision rule serves as the splitting criterion for growing the weighted regression tree. We do not use more sophisticated strategies such as pruning by cross validation, since boosting strong hypotheses may lead to overfitting. Empirically, GAMBLE performs reasonably well using regression trees with about 15 leaves.

3.2.2 GAMBLE with Weighted Regression Trees

In the general GAMBLE algorithm proposed earlier, two regression functions need to be fitted to the weighted samples in each iteration. We propose an implementation of the GAMBLE algorithm using weighted regression trees, which only requires one regression process. The key to this efficient solution lies in that \mathbf{y} only takes on categorical responses. Theorem 2 reveals the structural relationship of the two regression trees.

Theorem 2. *The topological structure of the weighted regression trees with respect to \mathbf{y} and \mathbf{z} ($z_j = y_j^2, j = 1, \dots, K$) is identical.*

Proof. See Appendix. □

Corollary 1. *The regression functions \mathbf{h} and \mathbf{g} satisfy*

$$\mathbf{h}(\mathbf{x}) = \frac{K-2}{K-1} \mathbf{g}(\mathbf{x}) + \frac{1}{K-1}$$

Proof. See Appendix. □

Theorem 2 implies that we only need to grow a regression tree T_1 with respect to \mathbf{y} , and the regression tree T_2 with respect to \mathbf{z} has exactly the same structure as T_1 . Formally, $\mathbf{h}(\mathbf{x})$ satisfies the same form as $\mathbf{g}(\mathbf{x})$,

$$\mathbf{h}(\mathbf{x}) = \sum_{j=1}^J \mathbf{d}_{R_j} I[\mathbf{x} \in R_j] \quad (3.9)$$

Algorithm 2 Gentle Adaptive Multiclass Boosting Learning (GAMBLE) Using Weighted Regression Trees.

Input: Training data set $S = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$

- 1: Initialization: set the observation weights $w_i \leftarrow \frac{1}{N}$, $i = 1, 2, \dots, N$, and set $\mathbf{F}(\mathbf{x}) \leftarrow 0$.
- 2: **for** $m \leftarrow 1$ to M **do**
- 3: Using weight distribution $\{w_i\}$, fit a weighted regression tree $\mathbf{g}^{(m)}(\mathbf{x})$ of \mathbf{y} to \mathbf{x} .
- 4: Set $\mathbf{r}^{(m)}(\mathbf{x}) = K \frac{(K-1)\mathbf{g}^{(m)}(\mathbf{x})}{(K-2)\mathbf{g}^{(m)}(\mathbf{x})+1}$.
- 5: Obtain weak learner $f_j^{(m)}(\mathbf{x}) = r_j^{(m)}(\mathbf{x}) - \frac{1}{K} \sum_{k=1}^K r_k^{(m)}(\mathbf{x})$, $j = 1, 2, \dots, K$.
- 6: Update committee $\mathbf{F}(\mathbf{x}) \leftarrow \mathbf{F}(\mathbf{x}) + \mathbf{f}^{(m)}(\mathbf{x})$.
- 7: Update weights $w_i \leftarrow w_i \exp(-\frac{1}{K} \mathbf{y}_i^T \mathbf{f}^{(m)}(\mathbf{x}_i))$, $i = 1, 2, \dots, N$ and renormalize.
- 8: **end for**

Output: the committee classifier: $C(\mathbf{x}) = \arg \max_k F_k(\mathbf{x})$.

Moreover, within the same region, Corollary 1 suggests that there is a simple way to transform the regression value from the regression function $g(\mathbf{x})$ to the (un-centralized) weak learner $r(\mathbf{x})$ (see Figure 3.1).

Theorem 2 and Corollary 1 clarify the relationship between the two regression functions based on the WLS criterion, thus yielding a more efficient solution (Algorithm 2). Obviously, Algorithm 2 works twice as fast as Algorithm 1 since it fits only one regression function in each step. In what follows, we refer to Algorithm 2 as GAMBLE unless stated otherwise.




Figure 3.1. The concave transformation of the regression function $g(x)$ to the (un-centralized) weak learner $r(x)$.

3.3 Remarks on the GAMBLE Algorithm

We characterize the relationship of the GAMBLE algorithm with other boosting variants, and remark on its superiority over other multiclass boosting alternatives.

First, GAMBLE is a natural extension of Gentle AdaBoost (see Algorithm 3), and they are identical up to a factor of 2 in the binary case. The difference comes from the different coefficients in the loss functions. In this regard, Gentle AdaBoost is the binary version of GAMBLE. Also, GAMBLE retains the succinct form of Gentle AdaBoost when generalizing to the multiclass case.

Algorithm 3 Gentle AdaBoost (Friedman et al. [8]).

- 1: Initialization: set the observation weights $w_i \leftarrow \frac{1}{N}$, $i = 1, 2, \dots, N$, and set $\mathbf{F}(\mathbf{x}) \leftarrow 0$.
 - 2: **for** $m \leftarrow 1$ to M **do**
 - 3: Using weights w_i , fit a regression function $g^{(m)}(\mathbf{x})$ to \mathbf{y} .
 - 4: Update $F(\mathbf{x}) \leftarrow F(\mathbf{x}) + g^{(m)}(\mathbf{x})$.
 - 5: Update weights $w_i \leftarrow w_i \exp(y_i \mathbf{g}^{(m)}(\mathbf{x}_i))$, $i = 1, 2, \dots, N$ and renormalize.
 - 6: **end for**
- Output:** Output the classifier: $C(\mathbf{x}) = \text{sign}[F(\mathbf{x})]$.
-

Second, the same argument about numerical stability for Gentle AdaBoost and AdaBoost.Real [8] holds when we compare GAMBLE and SAMME.Real. Specifically, we reexamine how GAMBLE obtains a weak classifier from a regression function in step (2b) (Algorithm 2). Result 2 and Result 3 imply that the regression functions and the weak classifiers are both bounded.

Result 2. *The regression function $\mathbf{g}(\mathbf{x})$ satisfies $\sum_{k=1}^K g_k^{(m)}(\mathbf{x}) = 0$; also, $\dim(\{\mathbf{g}^{(m)}\}) = K - 1$.*

Proof. See Appendix. □

Result 3. *For any \mathbf{x} and dimension $k \in \{1, 2, \dots, K\}$, $g_k^{(m)}(\mathbf{x}) \in [-\frac{1}{K-1}, 1]$ and $r_k^{(m)}(\mathbf{x}) \in [-K(K-1), K]$; the weak learner $f_k^{(m)}(\mathbf{x}) \in [-(K^2 - K), K^2 - K]$.*

Proof. The derivation follows from the proof of Theorem 2. See appendix for derivations. □

SAMME.Real, on the other hand, relies on the logit-transform of the posterior probability, estimated by the training samples within a region. Due to the scarcity of data, it is highly likely that no training samples of a particular class appear in this region. Thus the posterior probability is estimated to be zero. As such, the logit will be $-\infty$, making the weak learner explode. Even when smoothing methods are used to account for unseen data, the logit is still large, and the weight updates are exponentially larger. In later iterations, the weights of training samples differ by orders of magnitude and result in numerical instability. The larger the K , the more severe the problem will be. Thus it explains the numerical stability of GAMBLE compared to SAMME.Real.

Besides the numerical stability issue, the key difference between GAMBLE and SAMME is how the multiclass loss is optimized. Although the same loss function is used, GAMBLE adopts

adaptive quasi-Newton steps whereas SAMME.Real directly optimizes the Lagrange, the latter of which being more aggressive.

Finally, compared to Gentle AdaBoost.MH, GAMBLE is a natural extension for Gentle AdaBoost to the multiclass case. Moreover, GAMBLE is K times more efficient than Gentle AdaBoost.MH. Suppose a regression tree with depth d is used as the weak learner, the complexity of building it is $O(dDn \log(n))$, where D is the dimension of the input and n is the size of the training dataset. Note that GAMBLE builds only one regression tree per iteration, regardless of the number of fitted values. Thus its total cost is $O(MdDn \log(n))$, as compared to Gentle AdaBoost.MH's $O(KMdDn \log(n))$, where M is the number of iterations.

Query By Committee Style Active Learning

In practice, there are several learning settings where active learning is particularly valuable. First, labeled data is expensive to obtain due to the cost of human annotation, whereas unlabeled data is typically abundant. The cost of labeling can be minimized if the learner is able to learn from limited labeled data while actively querying only the most informative unlabeled instances. Second, if a tremendous amount of labeled data is available, boosting on the entire training dataset for even a small number of iterations can take large amount of time. Learning an optimal hypothesis by running a large number of iterations can be computationally infeasible. *Active learning* continuously augments the training dataset with the most informative instances at each step, thereby allowing the learner to focus on the ‘best’ portion of the training data. Thus it not only relieves the learner of redundant training samples, but it can also dispose of noisy samples.

In either case, Query By Committee (QBC) [19] is an effective active learning strategy. The

Algorithm 4 Active GAMBLE Algorithm.

- 1: Initialization: randomly choose a small set of samples from the query set $Q = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ to the working set S .
- 2: **for** $l \leftarrow 1$ to L **do**
- 3: Train a committee $C_l = \text{GAMBLE}(S)$.
- 4: Test C_l on data set $T = Q - S$.
- 5: For each $\mathbf{x} \in T$, compute utility $u(\mathbf{x}) = [F_k(\mathbf{x}) - F_j(\mathbf{x})]^{-1}$, where $k = \arg \max_k F_k(\mathbf{x})$ and $j = \arg \max_{j \neq k} F_j(\mathbf{x})$.
- 6: Let set D be the R samples $\mathbf{x} \in T$ with the highest utility, $Q \leftarrow Q - D$ and $S \leftarrow S \cup D$.
- 7: If Q is empty, let $C_L \leftarrow C_l$ and exit loop.
- 8: **end for**

Output: the classifier: C_L .

general QBC method reduces the error of the learner by choosing to label the sample that splits the version space into two parts of comparable size. As such QBC theoretically has near-optimal data efficiency, which is logarithmic in the probability of error [6]. Such a result in the general QBC framework is however based on the assumption of using Gibbs algorithm as the component learner, generally intractable in practice.

In Active GAMBLE (algorithm 4), we adopt the same QBC sample selection method as in QBoost [1], and extend it to the multiclass case. In each trial, Active GAMBLE trains a committee using GAMBLE as the component learner, based on the current working set rather than the entire training data. It then evaluates the potential utility of the unlabeled samples in the query set and selects the top R samples with the highest utility into the working set. *Utility* is here defined as the reciprocal of the difference between the committee’s output for the most and the second most popular class label. By Theorem 1, the committee’s output can be regarded as the surrogate of the posterior probability. As such utility measures the uncertainty in the committee about its predicted label, i.e. a sample has higher utility when the committee is less certain about its label. [17] has shown that the hypothesis having a larger margin generalizes better on unseen data. Since the definition of utility here is similar to the margin formulation in [17, 1] we expect that Active GAMBLE generalizes well by selecting samples with high utility or small margin, thereby increasing the margin in the hypothesis. Empirical results in Section 5.2 suggest that Active GAMBLE can potentially achieve higher classification accuracies. The results also show that Active GAMBLE greatly improves data efficiency and creates simpler model, therefore reduces the training and testing time. For very large training datasets, one can further reduce the training time of Active GAMBLE by only evaluating the utility of a random portion of instances in the query set.

Experiments

We conduct a series of experiments on both simulated and real world data to showcase the performance of the algorithm in various multiclass classification problems. We choose the SAMME.Discrete algorithm¹ (SAMME for short) in our experiments for comparison for several reasons. First, it has been shown that SAMME outperforms the most popular multiclass boosting algorithm AdaBoost.MH, in terms of both accuracy and efficiency [22]. In addition, GAMBLE and SAMME can be treated as variants because the same multiclass loss function and K-class response encoding scheme are used in both algorithms. Note that unlike SAMME which uses the Lagrange multiplier for direct optimization, GAMBLE uses the quasi-Newton steps to minimize the exponential loss. The goal is to provide a meaningful and efficient multiclass extension of Gentle AdaBoost, one of the best binary boosting algorithms.

Table 5.1. UCI benchmark datasets. K is the number of classes and D is the number of feature dimensions. The last two columns are the numbers of the training and testing instances, respectively.

Dataset	K	D	#Train	#Test
Glass	6	9	113	101
Pendigits	10	16	7,494	3,498
Satimage	6	36	4,435	2,000
Segmentation	7	19	210	2,100
Soybean	19	35	307	376
Splice	3	60	2,860	316
Vehicle	4	18	422	422
Vowel	11	10	528	462

¹SAMME.Discrete is numerically more stable than SAMME.Real, though empirical results have shown that they perform almost equally well [22].

5.1 Experiments on Real World Data

We tested our algorithm on a collection of benchmark datasets (Table 5.1) available from the UCI Machine Learning Repository² [14]. We intentionally choose the same multiclass datasets as in [18] so readers can compare our results with other multiclass boosting variants such as discrete and real AdaBoost.MH, AdaBoost.MR, etc. These real world datasets cover a wide variety of multiclass classification scenarios.




Figure 5.1. GAMBLE boosting regression trees of different sizes in the Pendigits dataset. T denotes the number of leaves in the regression tree. Note that the iteration number is shown on a log scale.

Figure 5.1 shows how GAMBLE performs when boosting regression trees of different sizes. Boosting stumps requires a lot more boosting steps than boosting sophisticated trees, in order to reach a desirable level of prediction error. Interestingly, the gap between test errors becomes more narrow when boosting relatively strong learners. For instance, the test errors of boosting trees with the number of leaves 10, 12 and 14 are almost the same after 600 iterations. Although boosting sophisticated learners may require fewer iterations for the convergence of the test error, each iteration takes longer training time and it may also risk overfitting the training data. Therefore we need to strike a balance in the size of the trees.

The test results of GAMBLE and SAMME in the UCI datasets are summarized in Figure 5.2. GAMBLE generally incurs less test error rates at 10, 100 and 1,000 iterations. The sign test reveals that the zero median hypothesis can be rejected ($p = 0.0066$) at the significance level

²<http://www.ics.uci.edu/~mlern/MLRepository.html>




Figure 5.2. Test error scatter plot for GAMBLE and SAMME at 10, 100 and 1,000 iterations in 8 UCI multiclass datasets. Each point shows the error rate of GAMBLE and SAMME on one benchmark dataset.

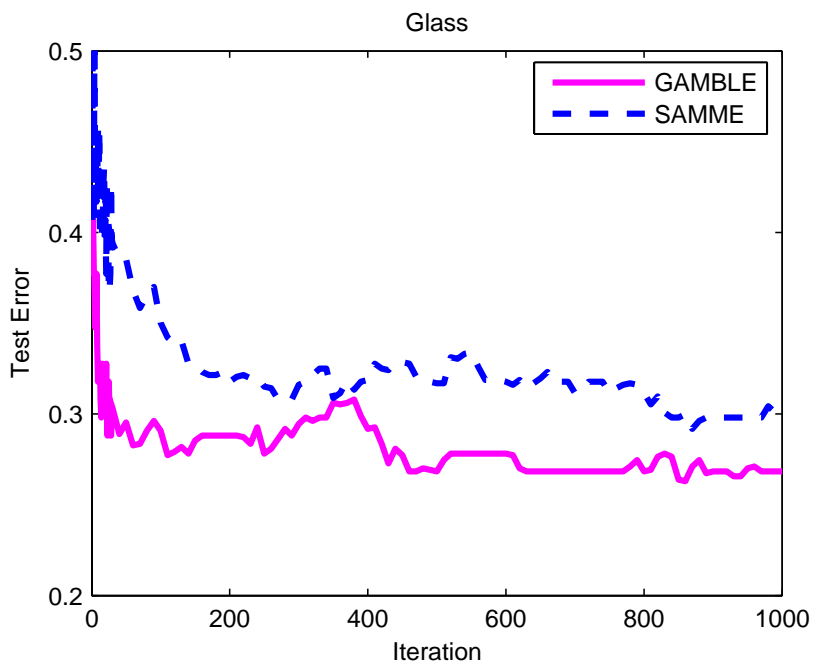


Figure 5.3. Test error (in the range of $[0, 1]$) for GAMBLE and SAMME. Discrete in the *Glass* dataset.

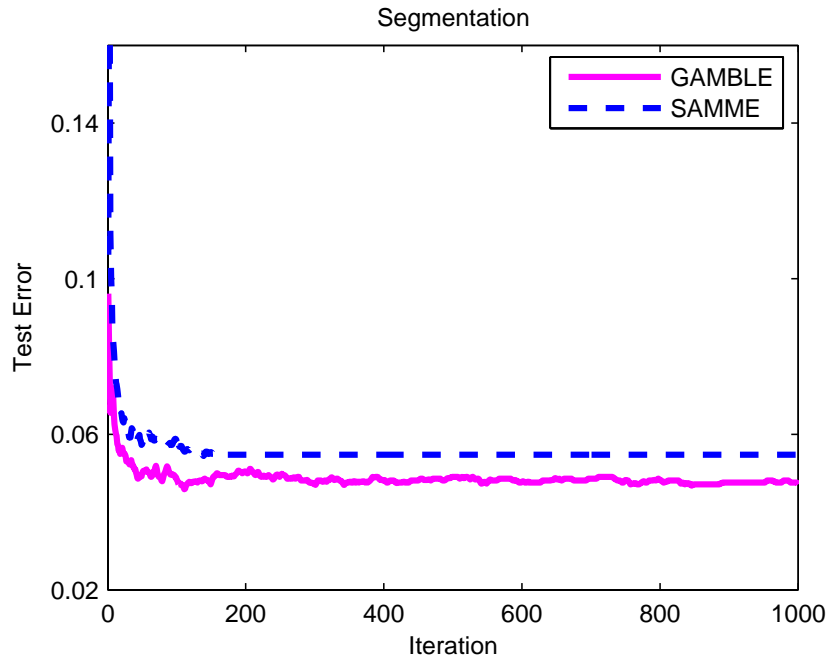


Figure 5.4. Test error (in the range of $[0, 1]$) for GAMBLE and SAMME.Discrete in the *Segmentation* dataset.

1%. In other words, GAMBLE performs statistically better than SAMME in these datasets. Figure 5.3 and Figure 5.4 illustrate the test errors of GAMBLE and SAMME in the *Glass* and *Segmentation* datasets, respectively. Both datasets, compared to their dimensions, have small number of training samples for each class. GAMBLE however can handle these sparse datasets and generalize well.

5.2 Active Learning Performance

We evaluate the effectiveness of the active learning strategy by test error, training time and data efficiency. *Data efficiency* is defined as the fraction of training samples used to achieve the target error rate. As a baseline algorithm, we implement the GAMBLE algorithm with random sampling, which we refer to as Random GAMBLE henceforth. Random GAMBLE selects training samples randomly from the training set regardless of their utility. The results of Random GAMBLE are averaged over 10 runs. The GAMBLE algorithm which uses the entire training dataset in each boosting step is referred to as Batch GAMBLE. For each dataset, the lowest error rate incurred by Batch GAMBLE throughout the first 1,000 iterations is set as the target error rate.

We present the results for two datasets which illustrate the characteristic learning curves of the active learning strategies. Figure 5.5 shows that in the *Pendigits* dataset, Active GAMBLE quickly reaches the target error rate (4.7%) using about 300 samples and plateaus afterwards.

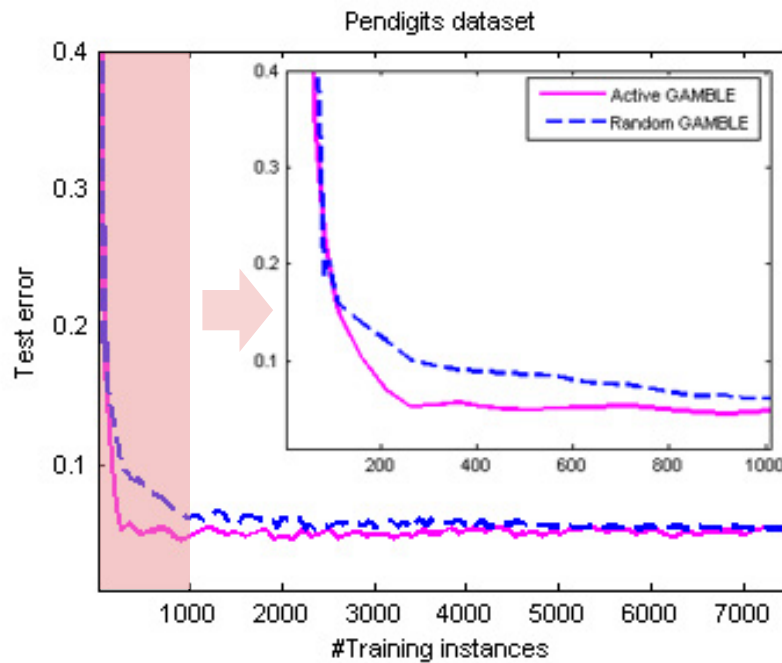


Figure 5.5. Test error of Active GAMBLE and Random GAMBLE in the *Pendigits* dataset. Active GAMBLE incurs lower test error and uses fewer training samples than Random GAMBLE. Note that the test error curves for the first 1000 samples are zoomed in for the *Pendigits* dataset.

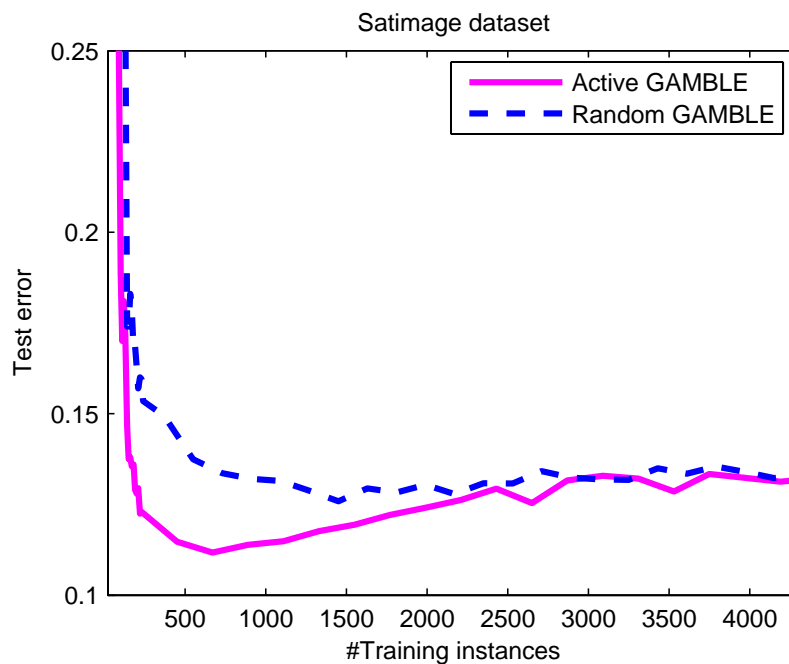


Figure 5.6. Test error of Active GAMBLE and Random GAMBLE in the *Satimage* dataset. Active GAMBLE incurs lower test error and uses fewer training samples than Random GAMBLE.

Table 5.2. Data efficiency comparison. *Lowest error rate* is the minimum error rate incurred by Active GAMBLE. *Target error rate* is the lowest error rate incurred by Batch GAMBLE in the first 1,000 iterations.

Dataset	Target Error Rate (%)	Lowest Error Rate (%)	# Training samples (pct.)		
			Total Samples	Active GAMBLE	Random GAMBLE
Glass	27.5%	27.5%	113	91 (80.5%)	102 (90.3%)
Pendigits	4.71%	3.92%	7,494	340 (4.39%)	2,510 (33.5%)
Satimage	12.6%	11.3%	4,435	220 (4.96%)	1,450 (32.7%)
Segmentation	4.62%	4.43%	210	135 (64.3%)	190 (90.5%)
Soybean	5.65%	5.25%	307	245 (79.8%)	300 (97.7%)
Splice	4.48%	3.42%	2,860	430 (15.0%)	970 (33.9%)
Vehicle	27.3%	25.1%	422	115 (27.3%)	225 (53.3%)
Vowel	47.0%	41.8%	528	155 (36.7%)	205 (48.6%)

Random GAMBLE converges to the same level of error rate using 1,000 samples. Figure 5.6 shows an even more interesting active learning curve. In the Satimage dataset, Active GAMBLE uses as few as 220 training samples (5% of all training data) to reach the target error rate (12.6%), while Random GAMBLE requires at least six times as many training samples. Interestingly, Active GAMBLE continues to reduce prediction error by about 2% after it reaches the target error rate. This implies that by using the 600 most informative samples (13.5% of all training data), Active GAMBLE is able to achieve higher classification accuracy. After this point, however, adding more training instances does not improve the accuracy and eventually the prediction error converges to that of Random GAMBLE.

Table 5.2 summarizes the results for the UCI datasets. By collecting the most informative instances, Active GAMBLE achieves far better **data efficiency** than Random GAMBLE, especially in the large training datasets, *Pendigits*, *Satimage* and *Splice*. This implies that those datasets contain redundant training samples. Batch GAMBLE requires global weight updates at each boosting step, whereas the component learner in Active GAMBLE only performs weight updates on a fraction of the training samples. Thus Active GAMBLE has faster **training time** than that of Batch GAMBLE, especially when high data efficiency is achieved (see Table 5.3). In

Table 5.3. Training time comparison.

Dataset	Training time (sec.)	
	Active GAMBLE	Batch GAMBLE
Glass	67	28
Pendigits	1,698	10,043
Satimage	362	6,066
Segmentation	158	72
Soybean	266	87
Splice	320	470
Vehicle	127	310
Vowel	136	367

small datasets such as *Glass*, *Segmentation* and *Soybean*, practically all training data is needed to converge to the target error rate. Batch GAMBLE works faster in these cases without having to search for the most informative samples. Table 5.2 also presents the lowest **test error rate** of Active GAMBLE, which in most cases is lower than the target error rate. In practice one can retain a portion of the training samples as a hold-out dataset to determine the early stopping point in training. This simpler model trained by less number of samples can then be used for prediction.

In conclusion, active learning strategy enables GAMBLE to work on large-scale training datasets, using only the portion of informative training samples and taking far less training time. Also, Active GAMBLE yields a simpler model which also reduces prediction time and potentially leads to better generalization.

5.3 Simulation Study

We tested GAMBLE with a popular synthetic example in boosting literature.

The Concentric Spheres Example

The concentric spheres example has been used in [8] to showcase the performance of different boosting algorithms. Since geometrically complex classification boundaries are involved in this example, it is regarded to be more complicated than practical classification problems. In this example, the samples are drawn from a ten-dimensional Gaussian distribution, i.e. $\mathbf{x} \sim N^{10}(\mathbf{0}, \mathbf{I})$. Each class is defined by thresholding the radius of the sphere $C_k = \{\mathbf{x}_i | t_{k-1} \leq r_i^2 \leq t_k\}$, where $r_i = \sqrt{\sum_{j=1}^{10} x_j^2}$ is the radius from the origin, $t_0 = 0$ and $t_k = \infty$. $\{t_k\}_{k=1}^K$ are chosen such that approximately equal number of instances are placed in each class. $K \cdot 1000$ instances are used as training samples, and an independently drawn set of 10,000 samples are used for testing.

Figure 5.7 represents the training and test error of GAMBLE when $K = 2$, and it reveals an impressive aspect of GAMBLE. Even after the training error reaches zero at about 20 iterations, GAMBLE continues to fit base classifiers to the training samples and the test error keeps dropping thereafter without indications of overfitting. This can be interpreted as GAMBLE increases the margin of the training samples [17] even when they are all correctly classified. Eventually this results in better generalization on the test data.

Figure 5.8 shows the training and test error curves of GAMBLE and SAMME. This figure demonstrates significant difference between the two algorithms. GAMBLE outperforms SAMME by incurring 10% less test error, which is consistent with the findings in [8]. Note that after 120 iterations, SAMME has lower training error whereas GAMBLE has substantially lower test error than SAMME. We can explain this by the inherent ‘gentle’ nature of the GAMBLE algorithm, which places less emphases on samples that are hard to learn. This small sacrifice in training error actually translates to better generalization power in prediction. But in this complicated classification problem, SAMME overemphasizes the outliers of the Gaussian distribution and




Figure 5.7. GAMBLE’s training and test errors in the concentric spheres example (K=2). Iteration number is shown on the log scale, highlighting the significant difference in training error iterations to those of testing.




Figure 5.8. GAMBLE’s training and test errors in the concentric spheres example (K=5). Iteration number is shown on the log scale, highlighting the significant difference in training error iterations to those of testing.

drives the weak learners to fit them. As such this affects the prediction performance. This further confirms explanations in Chapter 3.3.

This example elicits two important aspects of GAMBLE. GAMBLE continuously increases the margin regardless of the training error and is less sensitive to outliers. Both of these are crucial to prediction error reduction. In practice, classification problems are less noisy and less complex than this example and the difference in performance is not as much pronounced.

Chapter 6

Conclusions

This work extends the successful Gentle AdaBoost algorithm to the multiclass case. The proposed GAMBLE algorithm follows closely the rationale of other boosting predecessors. It fits one weighted weak learner symmetrically for all classes in each boosting step and combines them into a powerful committee classifier. Compared to other ‘*one-against-all*’ variants, it reduces the time complexity by K . GAMBLE is numerically stable since all the weak learners are bounded, and is thus more robust to unseen data. Also, the only parameter required in GAMBLE is the complexity of the weak learner (e.g. the number of leaves in the regression tree). As such it demands minimum domain knowledge about the classification problem. Our experiments demonstrate that GAMBLE does not seem to overfit the training data after a large number of boosting steps, even after the training error reaches zero. To accelerate learning in large-scale datasets, we use active learning to select the most informative samples for training. Compared to Batch GAMBLE, Active GAMBLE is able to achieve even higher classification accuracy, using only a portion of the training samples and consuming far less training time. Our empirical results indicate that the proposed GAMBLE algorithm outperforms the state-of-the-art multiclass boosting methods.

Proofs of Theorems

A.1 Proof of Theorem 2

Proof. For simplicity of notations, we provide the proof assuming that the regression function generates one response, and it is relatively straightforward to generalize to the K-response case. Suppose during the regression process with respect to \mathbf{y} , the region R in question is split into two subregions R_1 and R_2 ($R_1 \cup R_2 = R$), subject to the criterion (3.8). In other words, the following inequality holds for any two subregions S_1 and S_2 ($S_1 \cup S_2 = R$):

$$\begin{aligned}
 & \sum_{\mathbf{x}_l \in R_1} w_l (y_l - c_{R_1})^2 + \sum_{\mathbf{x}_l \in R_2} w_l (y_l - c_{R_2})^2 \\
 & \leq \sum_{\mathbf{x}_l \in S_1} w_l (y_l - c_{S_1})^2 + \sum_{\mathbf{x}_l \in S_2} w_l (y_l - c_{S_2})^2
 \end{aligned} \tag{A.1}$$

Substituting (3.7) into the weighted squares and canceling out the identical weighted square response $\sum_{\mathbf{x} \in R} w_l y_l^2$ on both sides, (A.1) becomes:

$$\begin{aligned}
 & \frac{\left(\sum_{\mathbf{x}_l \in R_1} w_l y_l \right)^2}{\sum_{\mathbf{x}_l \in R_1} w_l} + \frac{\left(\sum_{\mathbf{x}_l \in R_2} w_l y_l \right)^2}{\sum_{\mathbf{x}_l \in R_2} w_l} \\
 & \geq \frac{\left(\sum_{\mathbf{x}_l \in S_1} w_l y_l \right)^2}{\sum_{\mathbf{x}_l \in S_1} w_l} + \frac{\left(\sum_{\mathbf{x}_l \in S_2} w_l y_l \right)^2}{\sum_{\mathbf{x}_l \in S_2} w_l}
 \end{aligned} \tag{A.2}$$

Since $y_l \in \{1, -\frac{1}{K-1}\}$, we further partition R_1 into R_1^+ when $y_l = 1$ and otherwise R_1^-

$(R_1^+ \cup R_1^- = R_1)$. Hence,

$$\begin{aligned} \sum_{x_l \in R_1} w_l y_l &= \sum_{x_l \in R_1^+} w_l \cdot 1 + \sum_{x_l \in R_1^-} w_l \cdot \left(-\frac{1}{K-1}\right) \\ &= \sum_{x_l \in R_1} w_l - \frac{K}{K-1} \sum_{x_l \in R_1^-} w_l \end{aligned} \quad (\text{A.3})$$

Applying the same trick to all the components and canceling out the common terms on both sides of (A.2), the following inequality holds:

$$\begin{aligned} &\frac{\left(\sum_{x_l \in R_1^-} w_l\right)^2}{\sum_{x_l \in R_1} w_l} + \frac{\left(\sum_{x_l \in R_2^-} w_l\right)^2}{\sum_{x_l \in R_2} w_l} \\ &\geq \frac{\left(\sum_{x_l \in S_1^-} w_l\right)^2}{\sum_{x_l \in S_1} w_l} + \frac{\left(\sum_{x_l \in S_2^-} w_l\right)^2}{\sum_{x_l \in S_2} w_l} \end{aligned} \quad (\text{A.4})$$

Note that $z_l = y_l^2 \in \{1, \frac{1}{(K-1)^2}\}$, therefore,

$$\begin{aligned} \sum_{x_l \in R_1} w_l z_l &= \sum_{x_l \in R_1^+} w_l \cdot 1 + \sum_{x_l \in R_1^-} w_l \cdot \frac{1}{(K-1)^2} \\ &= \sum_{x_l \in R_1} w_l - \frac{K^2 - 2K}{(K-1)^2} \sum_{x_l \in R_1^-} w_l \end{aligned} \quad (\text{A.5})$$

Also note that

$$\begin{aligned} &\frac{\left(\sum_{x_l \in R_1} w_l z_l\right)^2}{\sum_{x_l \in R_1} w_l} + \frac{\left(\sum_{x_l \in R_2} w_l z_l\right)^2}{\sum_{x_l \in R_2} w_l} \\ &= \sum_{x_l \in R_1} w_l + \sum_{x_l \in R_2} w_l - 2 \frac{K^2 - 2K}{(K-1)^2} \left[\sum_{x_l \in R_1^-} w_l + \sum_{x_l \in R_2^-} w_l \right] \\ &\quad + \frac{(K^2 - 2K)^2}{(K-1)^4} \left[\frac{\left(\sum_{x_l \in R_1^-} w_l\right)^2}{\sum_{x_l \in R_1} w_l} + \frac{\left(\sum_{x_l \in R_2^-} w_l\right)^2}{\sum_{x_l \in R_2} w_l} \right] \end{aligned} \quad (\text{A.6})$$

Substituting (A.4), we derive this inequality:

$$\begin{aligned}
(A.6) &\geq \sum_{x_l \in R} w_l - 2 \frac{K^2 - 2K}{(K-1)^2} \sum_{x_l \in R^-} w_l \\
&+ \frac{(K^2 - 2K)^2}{(K-1)^4} \left[\frac{\left(\sum_{x_l \in T_1^-} w_l \right)^2}{\sum_{x_l \in T_1} w_l} + \frac{\left(\sum_{x_l \in T_2^-} w_l \right)^2}{\sum_{x_l \in T_2} w_l} \right] \\
&= \frac{\left(\sum_{x_l \in T_1} w_l z_l \right)^2}{\sum_{x_l \in T_1} w_l} + \frac{\left(\sum_{x_l \in T_2} w_l z_l \right)^2}{\sum_{x_l \in T_2} w_l}
\end{aligned}$$

Thus we arrive at the following inequality,

$$\begin{aligned}
&\sum_{\mathbf{x}_l \in R_1} w_l (z_l - d_{R_1})^2 + \sum_{\mathbf{x}_l \in R_2} w_l (z_l - d_{R_2})^2 \\
&\leq \sum_{\mathbf{x}_l \in S_1} w_l (z_l - d_{S_1})^2 + \sum_{\mathbf{x}_l \in S_2} w_l (z_l - d_{S_2})^2
\end{aligned} \tag{A.7}$$

where d is the weighted mean for the square response.

Hence the optimal splitting variable and value that minimize the weighted least squares with respect to \mathbf{y} also minimize that with respect to \mathbf{z} . Thus, we can reach the conclusion by induction. \square

A.2 Proof of Corollary 1

Proof. Substitute (A.3) into (3.7) and similarly (A.5) into the weighted mean of \mathbf{z} , we derive the following equation:

$$\mathbf{d}_R = \frac{1}{K-1} + \frac{K-2}{K-1} \mathbf{c}_R \tag{A.8}$$

By definition of $\mathbf{g}(\mathbf{x})$ and $\mathbf{h}(\mathbf{x})$,

$$\begin{aligned}
\mathbf{h}(\mathbf{x}) &= \frac{K-2}{K-1} \sum_{j=1}^J \mathbf{c}_{R_j} I[\mathbf{x} \in R_j] + \frac{1}{K-1} \sum_{j=1}^J I[\mathbf{x} \in R_j] \\
&= \frac{K-2}{K-1} \mathbf{g}(\mathbf{x}) + \frac{1}{K-1}.
\end{aligned}$$

\square

Proofs of Results

B.1 Proof of Result 2

Proof. In each region R_j ,

$$\sum_{k=1}^K g_k^{(m)}(\mathbf{x} \in R_j) = \sum_{k=1}^K C_{k,R_j} = \sum_{k=1}^K \frac{\sum_{\mathbf{x}_l \in R_j} w_l y_{k,l}}{\sum_{\mathbf{x}_l \in R_j} w_l} = \frac{\sum_{\mathbf{x}_l \in R_j} w_l \sum_{k=1}^K y_{k,l}}{\sum_{\mathbf{x}_l \in R_j} w_l} = 0 \quad (\text{B.1})$$

which follows from the fact that the encoding scheme imposes the sum of the K components of the true response to be zero. Thus the K components of $\mathbf{g}^{(m)}(\mathbf{x})$ are linear dependent. However, any $K - 1$ components are independent, otherwise it suggests redundant class labels. Therefore, the dimension of the vector space spanned by $\{g_k^{(m)}\}_{k=1}^K$ is $K - 1$. \square

B.2 Proof of Result 3

Proof. Substitute (A.3) into (3.7), \mathbf{c}_{R_j} can be expressed as:

$$c_{k,R_j} = 1 - \frac{\sum_{\mathbf{x}_l \in R_{k,1}^-} w_l}{K - 1 \sum_{\mathbf{x}_l \in R_{k,1}} w_l} \quad (\text{B.2})$$

Since all the weights are positive, $-\frac{1}{K-1} \leq c_{k,R_j} \leq 1$. Therefore, $g_k^{(m)}(\mathbf{x})$ lies in the range $[-\frac{1}{K-1}, 1]$. In this range, $r_k^{(m)}$ monotonically increases with $g_k^{(m)}$ (see Fig. 3.1), thus $r_k^{(m)}(\mathbf{x}) \in [-K(K-1), K]$. \square

Bibliography

- [1] ABE, N., AND MAMITSUKA, H. Query learning strategies using boosting and bagging. In *Proc. of the 15th International Conference on Machine Learning* (1998).
- [2] ALLWEIN, E., SCHAPIRE, R., AND SINGER, Y. Reducing multiclass to binary: A unifying approach for margin classifiers. *J. Machine Learning Res.* 1 (2000), 113–141.
- [3] BREIMAN, L., FRIEDMAN, J., OLSHEN, R., AND STONE, C. *Classification and Regression Trees*. Wadsworth, 1984.
- [4] DIETTERICH, T. An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization. *Journal of Machine Learning Research* 40, 2 (2000), 139–158.
- [5] FREUND, Y., AND SCHAPIRE, R. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences* 55, 1 (August 1997), 119–139.
- [6] FREUND, Y., SEUNG, H., SHAMIR, E., AND TISHBY, N. Selective sampling using the query by committee algorithm. *Machine Learning* 28 (1997), 133–168.
- [7] FRIEDMAN, J. Another approach to polychotomous classification. Tech. rep., Stanford Univ., 1996.
- [8] FRIEDMAN, J., HASTIE, T., AND TIBSHIRANI, R. Additive logistic regression: a statistical view of boosting. *The Annals of Statistics* 28 (2000), 337–407.
- [9] HASTIE, T., AND TIBSHIRANI, R. Classification by pairwise coupling. *The Annals of Statistics* 26 (1998), 451–471.
- [10] HASTIE, T., TIBSHIRANI, R., AND FRIEDMAN, J. *The Elements of Statistical Learning*. Springer-Verlag, 2001.
- [11] HUANG, J., ERTEKIN, S., SONG, Y., ZHA, H., AND GILES, C. L. Efficient multiclass boosting classification with active learning. In *Proceedings of The 7th SIAM International Conference on Data Mining (SDM)* (2007), pp. 297–308.
- [12] LIN, Y. A note on margin-based loss functions in classification. *Stats. and Prob. Letters* 68, 1 (2004), 73–82.
- [13] MEIR, R., AND RATSCH, G. An introduction to boosting and leveraging. *Adv. Lect. on Machine Learning* (2003).

- [14] MERZ, C., AND MURPHY, P. UCI repository of machine learning databases.
- [15] SCHAPIRE, R. Using output codes to boost multiclass learning problems. In *Proc. of the 14th International Conference on Machine Learning* (1997), Morgan Kaufman, San Francisco, pp. 313–321.
- [16] SCHAPIRE, R. The boosting approach to machine learning: an overview. In *MSRI Workshop on Nonlinear Estimation and Classification* (2002).
- [17] SCHAPIRE, R., FREUND, Y., BARLETT, P., AND LEE, W. S. Boosting the margin: A new explanation for the effectiveness of voting methods. In *Proc. of the 14th Int'l Conf. on Machine Learning (ICML)* (1997).
- [18] SCHAPIRE, R., AND SINGER, Y. Improved boosting algorithms using confidence-rated prediction. *Machine Learning* 37, 1 (1999), 297–336.
- [19] SEUNG, H. S., OPPER, M., AND SOMPOLINSKY, H. Query by committee. In *Proc. of the 5th Annual Workshop on Computational Learning Theory* (1992).
- [20] VAPNIK, V. *The Nature of Statistical Learning Theory*. Springer-Verlag, 1995.
- [21] VAPNIK, V. *Statistical Learning Theory*. John Wiley and Sons, Inc., New York, 1998.
- [22] ZHU, J., ROSSET, S., ZHOU, H., AND HASTIE, T. Multiclass adaboost. Tech. Rep. #430, Department of Statistics, University of Michigan, 2005.