ROMA TRE
UNIVERSITÀ DEGLI STUDI

FACOLTÀ DI ECONOMIA

SCUOLA DOTTORALE IN ECONOMIA E METODI QUANTITATIVI
DOTTORATO IN METODI STATISTICI PER L'ECONOMIA E L'IMPRESA

XXIII Ciclo

Supervised Classification Methods: from Cross Validation to Forward Search Approach

Sabina Mazza

Supervisors: Prof. Giorgio Alleva
Prof. Marco Riani

Coordinatore: Prof.ssa Julia Mortera

Contents

CHAPTER 1
1 Introduction ..... Pag. 3
1.1 Overview ..... Pag. 3
1.2 Goals of the Thesis ..... Pag. 3
1.3 Outline of Chapters ..... Pag. 6

CHAPTER 2
2 The Stacking Framework ..... Pag. 8
2.1 Introduction ..... Pag. 8
2.2 Stacked Generalisation ..... Pag. 8
2.3 Stacking ..... Pag. 10
2.3.1 Meta-level Data ..... Pag. 10
2.3.2 Meta-level Classifiers ..... Pag. 10
2.4 StackingC ..... Pag. 11
2.5 Related Work ..... Pag. 12
2.6 Discussion ..... Pag. 13

CHAPTER 3
3 Advances in the Stacking scheme .....
3.1 Introduction ..... Pag. 14
3.2 Traditional elements of Stacking ensemble method ..... Pag. 14
3.2.1 Base Classifiers ..... Pag. 14
3.2.2 Meta-classifiers ..... Pag. 15
3.2.3 Classifiers Assessment ..... Pag. 17
3.3 Experimental set up of our Stacking proposal ..... Pag. 18
3.3.1 Software ..... Pag. 19
3.3.2 Simulation study for 0-level input data ..... Pag. 20
3.3.3 Base Classifiers ..... Pag. 21
3.3.4 Meta-classifiers ..... Pag. 23
3.3.5 Evaluating and comparing classification methods ..... Pag. 24
3.4 Discussion ..... Pag. 25

CHAPTER 4
4 Experimental Results ..... Pag. 26

4.1 Introduction ……. Pag. 26
4.2 Simulated data ……... Pag. 26
4.3 Simulated contaminated data …… Pag. 30
4.4 Are the number and type of base classifiers important? …… Pag. 34
4.5 Application to real data. Electrodes Data ……... Pag. 37
4.5.1 Results for STA13 …… Pag. 39
4.5.2 Results for STA6 ……... Pag. 40
4.5.3 Results for STA3 ……... Pag. 42
4.6 Application to real data: The dataset off-the-book employment ..... Pag. 43
4.7 Discussion ..... Pag. 50

CHAPTER 5
5 Forward Search in multivariate analysis in the context of supervised classification methods ..... Pag. 52
5.1 Introduction …… Pag. 52
5.2 The Forward Search …….. Pag. 52
5.2.1 Starting the Search …… Pag. 52
5.2.2 Progressing in the Search ……... Pag. 53
5.2.3 Monitoring the Search ……... Pag. 53
5.3 Extending Forward Search to Stacking schemes ..... Pag. 53
5.4 Discussion …… Pag. 58

CHAPTER 6
6 Experimental results ..... Pag. 59
6.1 Introduction …….. Pag. 59
6.2 Simulated data ..... Pag. 59
6.3 Simulated contaminated data …… Pag. 63
6.3.1 Simulated contaminated data in the hypothesis of a non-balanced search and a specific covariance matrix for each group ..... Pag. 76

6.4 Application to real data ..... Pag. 78
6.4.1 Electrodes Data …… Pag. 78
6.4.2 muscular Dystrophy Data and Off-the-book Workforce Data ..... Pag. 82
6.5 Discussion ..... Pag. 86

CHAPTER 7
Conclusions and Future Developments ..... Pag. 87

Appendix to Chapter 3 ..... Pag. 92

Appendix to Chapter 4 ..... Pag. 97

Appendix to Chapter 6 ..... Pag. 99

Bibliography ..... Pag. 129

CHAPTER 1

Introduction

In this chapter, which is the introduction, first we describe the framework in which this thesis is set that is a combination of supervised classification methods, in particular the Stacking Technique. Then we explain the motivation, goals and purposes of the thesis and the tools and methods used to achieve them. We finally conclude with the outline of the subsequent chapters.

1.1. Overview

Among those elements which may influence the precision and stability of a classification method are the size and quality of the data set used for the estimation. Even slight modifications to the data set may lead to the construction of different models.
In order to satisfy the need for models that are more stable and more precise in their predictions, various methods have been proposed by the literature, based on the combination of models from the same class, among which: Bagging (Breiman, 1996), Boosting (Freund and Schapire,1996), Random Forest (Breiman 2001), and on others based on the combination of predictions deriving from different supervised classification methods.
This approach is also known as an *ensemble of classifiers* in the supervised classification task. The trend of studies in this direction that starts with *Stacked Generalization* (Wolpert, 1992) is particularly interesting, and is consolidated by the proposals offered by *Stacking* (Ting and Witten, 1999) and *Stacking C* (Seewald 2002), which tackle and overcome crucial problems previously unsolved in continuity with the original theory.
This class of models aims to combine the predictions coming from a set of different supervised classification algorithms (*base-level classifiers*) by means of a *meta-level classifier* in order to improve performances. The main idea behind Stacking is to use the predictions of the base classifiers as attributes in a new training set that keeps the original class labels, and then combine them.
The presence of outliers in the dataset, could also alter the structure of the classification model, and cause the generation of predictions that might not be reliable.

1.2 Goals of the thesis

The proposal consolidated in the stacking framework and the relative advances in the research on the elements that characterise this scheme for the combination of classifiers were the starting point for this thesis which intends to investigate this theme further. The idea is to explore in greater detail some aspects that seem to be less developed in the literature and could contribute to the introduction of further elements into the research, side by side with those critical elements already highlighted in the report.
Most of the research on ensembles of classifiers tends to demonstrate that Stacking can perform comparably to the best of the base classifiers as selected by cross-validation, if not better. It is to be hoped that we can expect that the final classifier produced by Stacking is able to achieve better

performances in terms of accuracy than the best level-0 classifier. Otherwise the computational onus created by the complexity of the procedure would not be justified.
This has motivated us to investigate empirically the performance of the Stacking technique, also in terms of stability and strength, solving the problem of the combination of supervised classified methods by using two different approaches: one may be defined as traditional and the other as innovative.
To this end, together with the approach that we will define as *traditional* and that is inserted into the framework of StackingC and uses the combination of different base classification methods that are constructed and evaluated via cross-validation, an extension of the *Forward Search* (Atkinson, Riani and Cerioli, 2004) is proposed, so as to have a *robust* approach to the same problem.
Forward Search is a methodological proposal which, apart from allowing anomalous values to be identified, also makes it possible to monitor in an iterative way the effect exerted by each unit on the model and on the quantities of interest in each step of the search.
The "philosophy" at the heart of the Forward Search approach is the creation of a dynamic data analysis process, compared to that of a "static" type, supplied by the traditional approach.

The research trend described has established the following objectives for this work:

1) Evaluation of the base-level and meta-level classifiers in terms of their accuracy when there are modifications in the size of the data set
2) Evaluation of the effects caused by the presence of anomalous values in the data set on the performances of the base-level and meta-level classifiers and their comparison using two different approaches:
    - Traditional (Cross Validation)
    - Innovative (Forward Search)

3) Evaluation of the results of the simulation studies carried out to establish whether, and to what extent, the combination of classifiers makes it possible to improve performances compared to the use of a single classifier.
4) Underlining the influence that single observations may have on each classifier's rule of decision.
5) Monitoring the stability of the allocation rule.

On what we might define as the traditional level, a Stacking scheme is proposed that has some differences compared to the well-known one, both in terms of characteristics that are already present and with regard to the introduction of innovative elements.
In particular, with regard to the assumption at the base of the theory that "even small changes to the training set may cause large changes in the classifier induced by a learning algorithm", that Breiman (1996) defined as "instability", referring to instable algorithms as classification trees and, taking into account that little has been discovered about the relationship between training dataset parameter settings and the performance of base classifiers and meta-classifiers, we chose not to use well-known input datasets, contained in databanks and extensively represented when dealing with problems of supervised classification. Since the topic of the choice of input data is, in our opinion, an important part of the study, we chose not to use well-known datasets but rather to carry out a wide-ranging simulation study that involved the generation of datasets with different characteristics

for the modification of the quality and size of the estimate data. Moreover, taking into consideration the effects that the presence of atypical observations might have on the model too, a great deal of space has been assigned to the contamination design of the datasets which has introduced anomalous values with varying degrees of intensity, level and typology to allow us to explore this particular dimension of the problem, covered very little by the literature, and to give any relative indications.
Among the main differences from the predominant literature we find the choice to build the Stacking procedure entirely in a Matlab environment. The whole Stacking scheme has therefore been implemented in Matlab and built in the form of an interlinked process which is begun by the generation of the level-0 data in the case of the simulation study, and which includes a complete homogenisation of the procedures relative to each of its phases in order to guarantee uniformity and therefore comparability of the outputs returned at every step.
This allowed us to create a flexible and powerful work tool, although it is sometimes a little heavy in computational terms.
Since the Stacking algorithm, although included in other software (such as the open source software WEKA), did not exist in Matlab, a code for creating a complex structure was completely implemented, which made possible:

- The organisation of a complex and extended experimental plan in order to carry out a wide-ranging simulation study
- The building of a set of base classifiers and the appropriate procedure of cross-validation for carrying out the fit, the assessment and the generation of predictions for the formation of the input dataset for the meta-learner
- The fit and the assessment of the meta-classifier and thus the procedures for the prediction combination and the homogeneous processing of the results with regard to the characteristics of each method
- The creation of suitable plots

It was necessary to carry out a process of homogenisation for each step of the procedures for all the classifiers, as we pointed out, which were chosen voluntarily with different characteristics, in order to obtain the same output that is indispensable for making the structure function and for the assessment and comparability of the results.
There have been some extensions and modifications to some algorithms compared to the implementation provided for in Matlab, respecting all the decision rules that preside over individual functioning. Using the various implementations and modifications of the default parameters will provide an indication for each classifier in Chapter 3 in the Section dedicated to their description.
The innovative nature of the proposal for this part of the thesis is to be found chiefly in the extention of the Forward Search approach to the Stacking scheme in order to build the procedure in a robust way and to monitor the effects that each observation, outlier or not, can exert on the model and on the quantities of interest.

As we have already seen in the part of the thesis dedicated to the traditional approach, the entire process of the proposed Stacking scheme was implemented in a Matlab environment in this case too.

We created a specific and suitable routine which performs Forward Search in multivariate analysis in the context of the combination of supervised classification methods, which is then inserted into the field of the FSDA toolbox, created for multivariate data analysis.
This procedure is, in any case, more complex than the other one, because as well as the construction of the Stacking scheme, there is also its inclusion into a Forward Search context to be considered, which in turn has been extended to the field of supervised classification. Thus the typical Forward Search procedures are extended to Stacking, procedures which refer to the choice of the best robust subset, the criteria for search progress and the creation of specific plots that support the monitoring of the quantity of interest graphically.
Starting from the construction, using robust methods, of a subset S(m) free from anomalous values, which represents the heart of the distribution, a dynamic implementation will be achieved, thus increasing the size of the robust sample selected with the introduction of one observation at a time. The choice of the new cardinality subset $m+1$ is founded upon the Mahalanobis distances calculated to the step m. The observations will be chosen with *the* $m+1$ with the smallest distances to form the new subset $S(m+1)$. The process is repeated at every step of the search and continues until $m = n$.
This entire procedure was repeated for every step of the process to as to provide us with a dynamic monitoring of the performances of the classifiers in terms of classification error compared to the units that form the subset and the units outside it, of the quantities of interest such as posterior probabilities and cross validation probabilities and, of course, of the maximum Mahalanobis distances of the units inside the subset as well as of the minimum ones of the units outside of the subset.
For both approaches, some applications of the proposed Stacking Scheme to datasets generated by means of the experimental design and also real data are carried out. The first dataset used is "Electrodes Data" (in Atkinson, Riani and Cerioli (2004). The second application of the Stacking Scheme to real data was carried out on a data sample taken from a dataset containing the results of the inspection surveys carried out by INPS (National Social Security Institute) on Italian companies in order to see if there was any off-the-book employment present.

1.3 Outline of Chapters

In the following Chapter 2 we present the main proposals for the Stacking framework, giving a great deal of space to Wolpert's and its main extensions.
In the first place the one provided by Ting and Witten which tackles and resolves crucial problems previously unsolved that Wolpert defined "black art":

- the choice of the type of attributes that represent meta-level input data for the combiner function. They propose using the outputs represented by the probability distributions that are derived from using a set of base-level classifiers as *level-1 data* instead of the predictions of single class values as proposed by Wolpert.
- the choice of a level-1 generaliser in order to obtain improved accuracy using the stacked generalization method. They recommend the use of MLR (Multi-response linear regression) as a meta- level classifier, as used by Breiman (1996) in a Stacked regression learning scheme, and by Le Blanc and Tibshirami (1993). They believe that MLR is the best level-1 combiner when compared with other learning algorithms.

Then, the Ensemble scheme proposed by Seewald is illustrated, *StackingC*, which is based on reducing the dimensionality of the level-1 dataset not considering the entire probability distribution associated with each classifier as in Ting and Witten, (1999), but rather the dataset composed only of probability vectors expressed by each k base-level classifier on the belonging of the unit to a defined class.
Other proposals are presented that deal above all with the choice of the meta classifier such as that by Merz (1999) which proposes SCANN. This uses the Stacking strategies together with correspondence analysis to detect any correlations between the predictions of base-level classifiers, and as the meta-level combiner a nearest neighbor method is applied to predict unseen examples. Those which envisage the use of different types of Meta decision trees as meta-classifiers or those such as the contribution of Dzeroski and Zenko (2004) who propose two extensions of Stacking, one using an extended set of meta-level features and the other using multi-response model trees to learn at the meta-level. Finally an interesting proposal from Reid and Grudic (2009) which demonstrates empirically that even with a linear combination function, regularisation is necessary to reduce overfitting and increase predictive accuracy and propose different kind of regularisations.
In Chapter 3 we describe the proposed Stacking scheme, with particular attention to the traditional components of the Stacking process, by indicating the main differences between the proposed and the more well-known one. The whole procedure implemented in Matlab, the simulation study for 0-level input data, the contamination design, the extensions and modifications and the parameter settings implemented for each classifier and Stacking, are illustrated.
In Chapter 4 the empirical results obtained from the application of the proposed Stacking scheme to datasets generated by means of the experimental design and also to real data are shown. In particular, in section 4.2 the results relative to the non-contaminated data are illustrated, in order to investigate the effects on the performance of the base classifiers and of Stacking in the presence of input datasets with different characteristics. In Section 4.3 the application was carried out on simulated and contaminated data to investigate whether the presence of outliers can affect the performances of the base classifiers and of Stacking. In Section 4.4 are illustrated the results obtained using three different Stacking variants with different base-level classifier subsets built on different datasets In Sections 4.5 and 4.6 there are applications to real data.
Chapter 5 is dedicated to the extention of the Forward Search approach to the Stacking scheme. A brief and general illustration of the main phases of the Forward Search process in a multivariate context is given at the beginning of this chapter, followed by a description of the procedure for constructing the Stacking scheme which is then inserted into a Forward Search framework and the numerous plots returned that are a fundamental tool for identifying any anomalous values, for exploring data and for monitoring the performances of the various base classifiers and the final classifier, as well as the behaviour of the decision rule that presides over the functioning of each algorithm and of Stacking.
The experimental plan mentioned was also used for the Forward Search approach, and, in Chapter 6 some applications of the Stacking scheme proposed in a Forward Search context to simulated data e and to real data are illustrated.
Chapter 7 concludes this thesis with a summary of the results and an outline of future developments.

CHAPTER 2

Stacking Framework

2.1 Introduction

In this chapter we first describe the Stacking framework and then we summarize the main results of several recent studies of the Stacking technique for the combination of supervised classification methods.
The trend of studies that starts with *Stacked Generalization* (Wolpert, 1992) is particularly interesting, and is consolidated by the proposals offered by *Stacking* (Ting and Witten, 1999) and *Stacking C* (Seewald 2002), which tackle and resolve crucial problems previously unsolved in continuity with the original theory.

2.2 Stacked Generalisation

The aim of this ensemble learning scheme, originally proposed by Wolpert (1992), is to combine the predictions coming from a set of different supervised classification algorithms (*level-0 models*) by means of a *meta-level classifier* in order to improve prediction accuracy (as opposed to learning accuracy) as much as possible.
Test instance is first classified by each of the base classifiers. These classifications are fed into a meta-level training set from which a meta-classifier is produced.
The predictions of *level-0 classifiers* represent the attributes in a new training set (*level-1 data*), which keeps the original class labels. Stacking thus utilizes a *meta-learner* (*level-1 model*) to combine the predictions from different base classifiers which were estimated via cross-validation on a single data set.
There follows a brief description of the logic and the functioning of the Stacking technique together with a diagram (figure 1) which take into account some of the considerations made by Ting and Witten (1999) on the Wolpert proposal.
Given a set of K learning algorithms, called *level-0 generalisers* by Wolpert, and a data set :

$$L = \{(y_n, x_n), n = 1, ..., N\},$$

where y_n is the target value and x_n is a vector whose elements represent the values assumed by the variables for the *n-th* instance.

Let L be randomly split into J roughly equal-sized parts: $L_1, L_2, ..., L_J$

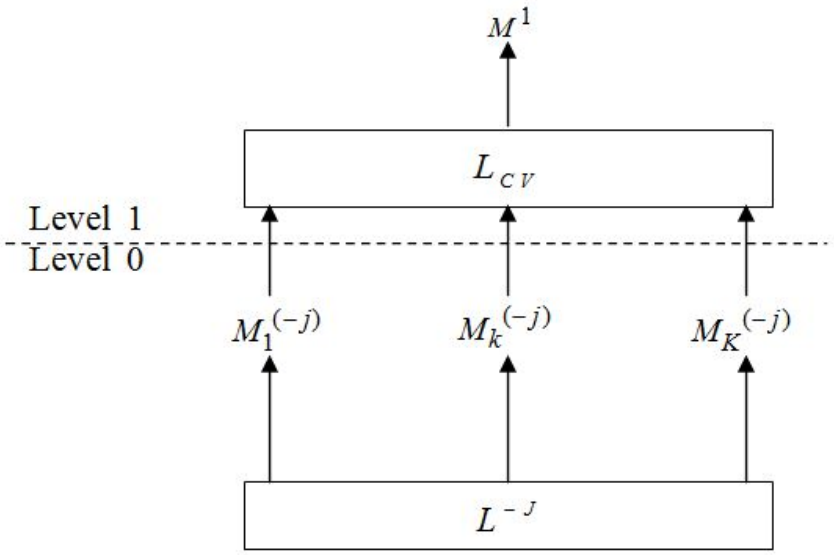

Figure 2.2.1: This figure illustrates the j-fold cross-validation process in level-0; the level-1 data set L_{CV} at the end of this process is used to produce the level-1 model M^1 (Ting and Witten, 1999).

We define :

L_j and $L^{(-j)} = L - L_j$ as the test and training set for the *j-th* fold of *J-fold* Cross Validation and, $M_k^{(-j)}$ a model for *k = 1,…K* is induced on the training set $L^{(-j)}$. Level-0 models.

For each vector x_n belonging to L_j, the test set for the *jth* cross-validation fold, let z_{k_n} be the prediction of $M_k^{(-j)}$ on x_n.

At the end of the cross-validation procedure, the dataset made up of the predictions of each K model using the terminology just introduced represents the level-1 data and is given by:

$$L_{CV} = \left\{ (y_n, z_{1n}, ..., z_{Kn}), n = 1, ..., N \right\}$$

The combiner function (*level-1 generaliser*) is then trained on this meta-level dataset to derive a Model M^1 (*level-1 model*) for y as a function of the predictions $(z_1, ..., z_k)$, whose output is the final classification of the units belonging to the input vector.

Formally, the final prediction function of *Stacked generalization* can be expressed by:

$$v_z(x) = \Sigma \alpha_k cz_z^k(x) \text{ for } k = 1, 2, ..., K .$$

Where cz_z^k is a set of *k* predictors.

This is the model proposed by Wolpert (1992) and universally considered to be the base model of stacked generalisation. It has been revisited and studied in depth by several scientists such as Breiman (1996) who demonstrated the success of stacked generalization in the setting of ordinary regression and Le Blanc and Tibshirami (1993)

However, it is interesting to note that Wolpert himself believes that many aspects of stacked generalization are, at present, a kind of "black art", and, therefore have not yet been resolved.

These aspects will be dealt with and resolved subsequently in continuity with the original theory, as we will see in the following sections.

2.3 Stacking

Ting and Witten (1999) with their Stacking learning scheme shed light on the following aspects that Wolpert himself believed to be a kind of "black art" in Stacked generalisation:

- the type of attributes that should be used to form level-1 data,
- the type of level-1 generaliser in order to obtain improved accuracy using the stacked generalization method.

2.3.1 Meta level data

Ting and Witten (1999) have proposed settings for the meta classifier and the type of meta data to be used in the field of the problems of supervised classification such as the extension of the application of Stacked Generalization .
They propose using the outputs represented by the probability distributions that are derived from a set of base-level classifiers as *level-1 data* instead of the predictions of single class values as proposed by Wolpert.
When returning to the notation and to the reference scheme already used to describe Stacked generalization, if a generic model $M_k^{(-j)}$ is used to classify an instance *x* belonging to L_j, and $P_{ki}(x)$ is the *probability* that *x* belongs to the *i-th* class, the following vector:

$$\mathrm{P}_{kn} = (P_{k1}(x_n),\ldots,P_{ki}(x_n),\ldots,P_{kI}(x_n))$$

represents the probabilities that the vector x_n belongs to the classes 1,..,*I*. assuming that classes have been returned by a single base-level classifier.
This gives the probabilities that the vector x_n belongs to the class 1,..,*I*. assuming that there are I classes and a set of k models with different bases.
The level 1 dataset will be composed of the aggregation of the probability vectors generated by all k models:

$$L'_{CV2} = \{(y_n, \mathrm{P}_{1n},...\mathrm{P}_{kn},...,\mathrm{P}_{Kn}), n = 1,..,N\}$$

Compared to the previous ensemble scheme of Stacked generalization, the final new model will be M^2.

2.3.2 Meta-level classifier

Ting and Witten propose the use of MLR (Multi-response linear regression) as a meta- level classifier, as used by Breiman (1996) in a Stacked regression learning scheme, and by Le Blanc and Tibshirami (1993). They believe that MLR is the best level-1 combiner when compared with other learning algorithms; it can represent a valid starting point in the search for the best method for meta-level learning to be used for problems in combining the supervised classification methods such as Stacking. Linear regression can easily be used for classification in domains with numeric attributes. Indeed, *any* regression technique, linear or nonlinear, is suitable for classification.

MLR is an adaptation of a least squares linear regression. For a classification problem with I class values, *I* separated regression problems are fitted: for each class *I*, a linear equation *LRI* is constructed to predict a binary response variable, which has value one if the class value is *I* and zero otherwise. Given a new example *x* to classify, *LRI* (*x*) is calculated for all *j* , and the class *k* is predicted with maximum *LRk* (*x*).
MLR, therefore, learns a linear regression function for each class which predicts a degree of confidence in class membership and can, after normalization, be interpreted as class probability. The output of the linear models, therefore, will have to be renormalized to yield a proper class probability distribution because the membership values they produce are not proper probabilities as they can fall outside the range 0 1.
Both Breiman (1996a) and LeBlanc & Tibshirani (1993) use the stacked generalization method in a regression setting and report that it is necessary to constrain the regression coefficients to be non-negative in order to guarantee that stacked regression improves predictive accuracy.

By modifying and simplifying Wolpert's hypothesis of Stacked generalization, seen in section 2.2, with regard to the final predictor: $v(x) = \Sigma \alpha_k v_k(x)$, the authors underline the need to enforce the non negativity of the coefficients $\{\alpha_k\}$, considering the hypothesis that the different v_k, by making predictions about the same data, could be strongly correlated and there may be no guarantee that the final (stacked) predictor is near the range which might degrade the *generalisation* performance of this learning method.
Ting and Witten (1999) have shown that non-negativity constraints on coefficients are not necessary.

2.4 StackingC

Seewald (2002) proposed an extension of Stacking, called StackingC, based on reducing the dimensionality of the level-1 dataset independently of the number of classes and removing a priori irrelevant features. In order to overcome a weakness of Stacking (Ting and Witten, 1999) in problems with more than two classes. StackingC seems to display better performances in terms of accuracy compared to Stacking, especially for multi-class problems, while for two-class datasets the improvements are more moderate, while the reduction of the size of the features makes a gain in computational terms.
The proposed method includes the use as input for the level-1 classifier (each linear model is associated with each of the classes), not the entire probability distribution associated with each classifier as in Ting and Witten, (1999), but rather the dataset composed only of probability vectors expressed by each k base-level classifier on the belonging of the unit to a defined class (Figure 2.4.1). In the learning scheme StackingC, therefore, each linear model learns as input data only those partial class probabilities that it is trying to predict.
The author maintains that the probability given by one classifier for only one class can be sufficient to guarantee the information necessary and also to ensure a good performance, because the sum of each class probability distribution has to be one, the probability of one class is one minus the probability of the other class.

Classifier 1 a	Classifier 2 a	…	Classifier K a	Class = a
$P1_a(x_1)$	$P2_a(x_1)$		$PK_a(x_1)$	1
$P1_a(x_2)$	$P2_a(x_2)$		$PK_a(x_2)$	0
.	.		.	.
.	.		.	.
$P1_a(x_N)$	$P2_a(x_N)$		$PK_a(x_N)$	0

Figure. 2.4.1. Level-1 data consisting only of partial probabilities given by each base-level classifier for class=a , k level-0 classifiers and N instances processed on the basis of the pattern proposed by Seewald (2002)

The use of MLR (Multi-Response Linear Regression) as a meta-level classifier is confirmed. Seewald tries to use other combiner functions instead of MLR, such as LWR (Locally Weighted Regression) and MP5Prime, a model tree learner implemented in the WEKA open-source software (Waikato Environment for Knowledge Analysis) developed at the University of Waikato in New Zealand. Empirically he finds that for two-class datasets MLR is the best classifier, even if the differences are minimal.
The author believes that, in this case, the source of the improvement lies partially in the dimensionality reduction, but more importantly in the higher diversity of class models that are combined.

2.5 Related Work

There have been several studies on combining classification models, including of course those on the Stacking framework.
The purpose of most of this research has been to study in depth those aspects defined by Wolpert as "black art" and therefore a great deal of attention has been paid to the choices in terms of meta data and meta-level classifiers.
There are several interesting proposals and the main ones will be looked at in brief below.
Merz (1999) proposes a method called SCANN (Stacking Correspondence Analysis and Nearest Neighbour) that uses the strategies of Stacking and correspondence analysis detect any correlations between the predictions of base-level classifiers, because it is well known that the combination of different classifiers improves the accuracy performance if they are weakly correlated. The original meta-level feature space (the class-value predictions) is transformed into a space of uncorrelated features. As the meta-level combiner a nearest neighbor method is applied to predict unseen examples. The author compares SCANN with two other stacking schemes that have a Naïve Bayes classifier as a meta-learner and a back-propagation trained neural network. Merz applied SCANN in this work to classifiers that only return class value predictions and not class probability distributions as in Stacking.
Todorovski and Dzeroski (2000) introduced a new algorithm to be used as a level-1 classifier: the meta decision Trees (MTDs), whose leaves do not contain class value predictions. Instead the most appropriate base level classifier to be used for classifying the unit that falls in that leaf is indicated. As first level dataset attributes they do not propose the use of probability distributions, but rather their characteristics, such as entropy and maximum probability, since they may be interpreted as estimates of the confidence of the classifier in its predictions.

Zenko et al. (2001) report that MDTs perform slightly worse compared to stacking with MLR. Overall, SCANN, MDTs, stacking with MLR and SelectBest seem to perform at about the same level. It would seem natural to expect that ensembles of classifiers induced by stacking would perform better than the best individual base-level classifier: otherwise the extra work of learning a meta-level classifier does not seem justified. The experimental results, however, do not show clear evidence of this.
Todorovski and Dzeroski (2002) report that stacking with MDTs makes it possible to exploit better than voting the differences between the base-level classifiers and has a better performance, especially in the hypothesis in which the mistakes made by the base level classifiers are uncorrelated. It is also superior when compared with SCANN, and the main ensemble methods of weak learners (especially decision trees) such as bagging and boosting.
Dzeroski and Zenko (2004) propose two stacking extensions with MLR, one using an extended set of meta-level features and the other using multi-response model trees instead of MLR as meta-classifiers. Firstly, the authors use the probabilities predicted for each class by each base classifier as meta-level features (as proposed by Ting and Witten) but augment with two additional sets of meta-level attributes: the probability distributions multiplied by the maximum probability and the entropies of the probability distributions. The results of their experiments show that there are no significant improvements when using only these two attributes (without the probability distributions), but when using all three sets of features at the same time, some improvements are noticeable. The second extension considers an alternative for MLR as meta-classifier, introducing Stacking with multi-response model trees, because model trees have been shown to perform better than MLR for classification via regression.
Reid and Grudic (2009) return to the need to insert constraints on coefficients; in fact they demonstrate empirically that with a linear combination function, regularization is necessary in order to improve accuracy and reduce overfitting. They propose using Ridge regression, lasso regression and elastic net regression because Stacking has a tendency to overfit, especially when highly correlated and well-tuned combination methods are used.

2.6 Discussion

We have outlined the main proposals of the literature that examine in depth and extend the Stacking Technique with particular attention paid to the choice of meta data and meta classifiers. In the thesis, as indicated, we will use the Stacking C ensemble scheme as a starting point for our analysis, but we plan to focus our attention on an exploration of the parameters, as well as the choice of meta data and meta classifiers. We also focus on an aspect that has been covered much less by the studies and that in our view, however deserves special attention: the choice of the initial dataset. This is connected to the assumption that small changes in the dataset can lead to different models and that the presence of outliers might alter the parameters of the model.

CHAPTER 3

Advances in the Stacking scheme

3.1 Introduction

The proposals illustrated in the previous chapter for the Stacking framework and the relative progress made in the research on the elements that characterize such a scheme of classifier combination are a valid starting point for this thesis, which intends to investigate this topic further. The idea is to explore in more detail some of the aspects that seem less developed and could contribute to the introduction of further elements into the research, side by side with the critical elements already highlighted in this work.
In the following section some of the components of the Stacking technique will be explained, especially in the usual outlook, while in section 3.3, the elements that are the essential aspects for constructing our Stacking model will be introduced, with clarifications regarding the main differences compared to the traditional formulation, both in terms of modifications in the characteristics of elements already found and with regard to the introduction of innovative elements.

3.2 Traditional elements of Stacking ensemble method

As we have demonstrated several times, elements traditionally considered to be critical for dealing with problems in the combination of supervised classification methods, and in particular in the ensemble Stacking method, are represented by the choice of base classifiers, meta classifiers and also by the type of meta data to be used. In our opinion, another important aspect is the assessment of the classifiers that will be illustrated in the section 3.2.3.
These elements will be summarized below, but it should be clear that there will not be a thorough examination of this theme, since many other types of learning algorithms could be used to generate base classifiers, and other typologies of meta-classifiers, used to provide a final prediction, but usable for describing the components that will be inserted in the Stacking process built in this thesis and described in section 3.3.

3.2.1 Base Classifiers

The base classifiers that will be used to build the proposed Stacking scheme, are methods having different characteristics because the learning algorithms that generated them are different. We made a voluntary choice to use classifiers in the combination that have different predictive capacities and strengths together with different decision rules for investigating whether the combination is able to enhance the performances of the most capable and mitigate the weaknesses of the less able

performers, and therefore Stacking can perform comparably to the best of the individual classifiers, if not better.

To make this treatment easier we can distinguish three categories among the algorithms that we will use such base classifiers in to the experimental set up:

- **Parametric methods**
  - Linear Discriminant Analysis
  - Quadratic Discriminant Analysis
  - Logistic Regression
  - Naive Bayes

- **Non-Parametric Methods**
  - Classification Tree
  - Support Vector Machine

- **Ensemble Methods**
  - Bagged Classification Tree
  - AdaBoost

The general formulations of the proposed algorithms will be summarised in Appendix, while in section 3.3.3 specific implementations carried out in a Matlab environment and relative to each algorithm will be illustrated.

3.2.2 Meta-classifiers

The most interesting of the proposed meta-classifiers are the following:

- Multi Response Linear regression (MLR)

- Ridge Regression

Multi-response Linear Regression is an adaptation of a least squares linear regression recommended (Ting and Witten 1999) for meta-level learning while several learning algorithms are shown not to be suitable for this task.
For a classification problem with K class values, *K* separated regression problems are fitted: for each class *k* , a linear equation *LR*k is constructed to predict a binary response variable, which has value 1 if the class value is *k* , and 0 otherwise. Given a new example *x* to classify, *LRk* (*x*) is calculated for all *j* , and the class *k* is predicted with maximum *LRk* (*x*). MLR, therefore, learns a linear regression function for each class which predicts a degree of confidence in class membership and can, after normalisation, be interpreted as class probability. The output of the linear models, therefore, will have to be renormalized to yield a proper class probability distribution because the membership values it produces are not proper probabilities as they can fall outside the range 0-1.

By using the cross-validated predictions $\hat{f}^{-k}(x)$ at *x*, using model *m*, applied to the dataset with the *i*th training observation removed. The stacking estimate of the weights is obtained from the least squares linear regression of y_i *on* $\hat{f}_m^{-1}(x)$, *m=1,2,*.....M.
The stacking weights are given by:

$$\hat{w}^{st} = \arg\min_{w} \sum_{i=1}^{N} \left[y_i - \sum_{m=1}^{M} w_m \hat{f}_m^{-i}(x_i) \right]^2 .$$

The final prediction then is $\sum_m \hat{w}_m^{st} \hat{f}_m(x)$.
Hastie et al. (2008) believe that better results can be obtained by restricting the weights to be nonnegative, and to sum to 1. This seems like a reasonable restriction if we interpret the weights as posterior model probabilities.

Ridge Regression

Ridge Regression, introduced by Hoerl and Kennard (1970), shrinks the regression coefficients by imposing a penalty on their size. The ridge coefficients minimize a penalized residual sum of squares

$$\hat{\beta}^{ridge} = \arg\min_{\beta} \left\{ \sum_{i=1}^{N} \left(y_i - \beta_0 - \sum_{j=1}^{p} x_{ij} \beta j \right)^2 + \lambda \sum_{j=1}^{p} \beta_j^2 \right\} \qquad (1.1)$$

Here $\lambda \geq 0$ is a complexity parameter that controls the amount of shrinkage: the larger the value of λ, the greater the amount of shrinkage. The coefficients are shrunk towards zero (and each other).
An equivalent way to write the ridge problem is:

$$\hat{\beta}^{ridge} = \arg\min_{\beta} \sum_{i=1}^{N} \left(y_i - \beta_0 - \sum_{j=1}^{p} x_{ij} \beta j \right)^2 \text{ subject to } \sum_{j=1}^{p} \beta_j^2 \leq t \qquad (1.2)$$

which makes explicit the size constraint on the parameters. There is a one-to-one correspondence between the parameters λ in (1.1) and *t* in (1.2). When there are many correlated variables in a linear regression model, their coefficients can become poorly determined and exhibit high variance. By imposing a size constraint on the coefficients, as in (1.2), this problem is alleviated. the intercept β_0 has been left out of the penalty term. The solution to (1.1) can be separated into two parts, after reparametrization using *centered* inputs: each x_{ij} gets replaced by $x_{ij} - \bar{x}_j$.

We estimate β_0 by $\hat{y} = \frac{1}{N} \sum_1^N y_i$.

The other coefficients get estimated by a ridge regression without intercept, using the centered x_{ij} .
Henceforth we assume that this centering has been done, so that the input matrix X has *p* (rather than *p* + 1) columns.

Writing the criterion in (1.1) in matrix form,

$$RSS(\lambda) = (y - X\beta)^T (y - X\beta) + \lambda\beta^T\beta$$

the ridge regression solutions are easily seen to be:

$$\beta^{ridge} = \left(X^T X + \lambda I\right)^{-1} X^T y$$

where I is the $p \times p$ identity matrix. Notice that with the choice of quadratic penalty $\beta^T\beta$, the ridge regression solution is again a linear function of y. The solution adds a positive constant to the diagonal of $X^T X$ before inversion. This makes the problem nonsingular, even if $X^T X$ is not of full rank, and was the main motivation for ridge regression when it was first introduced in statistics (Hastie et al. 2008). Ridge Regression is recommended such a meta-combiner in a Stacking scheme by Le Blanc and Tibishirami (1993), Breiman (1996) and, recently, Reid and Grudic (2009).

3.2.3 Classifiers Assessment

The generalization performance of a learning method relates to its prediction capability on independent test data. In classification task, we are interested to assess the ability of a learning algorithm to generalize on unseen data. It is common to measure a classifier's performance in terms of accuracy. Where
Accuracy = 1 - generalization error rate

It is our choice to measure and compare the performances of the classifiers based on their prediction error rate. The error rate is the proportion of misclassified instances over a whole set of instances, and it measures the overall performance of the classifier.
Of course one can be interested in the likely future performance on new data, because the error rate on the training set is not likely to be a good indicator of future performance.

$$Tra_{err} = \frac{1}{N}\sum_{i=1}^{N} L\left(y_i, \hat{f}(x_i)\right)$$

Any estimate of performance based on that data will be optimistic. Training error consistently decreases with model complexity, typically dropping to zero if we increase the model complexity sufficiently. However, a model with zero training error is overfitted to the training data and will typically generalize poorly. (Hastie et al. 2009).
Test error, or generalization error, is the prediction error on an independent test sample given by a classification method $\hat{f}(x)$ that has been estimated from a training set. To predict the performance of a classifier on new instances, we need to evaluate its generalization error rate on a dataset that has not been part of the classifier's fit. The test data must not be used in any way to build the classifier.
When the amount of data for splitting in training and test set is limited, one of the simplest and most popular methods for estimating prediction error is *K*-fold cross-validation.

We first split the data into *K roughly* equal parts. Then for each $k = 1, \ldots, K$, we remove the *kth* part from our data set and fit a model $\hat{f}^{-k}(x)$.
Let *Ck* be the indices of observations in the *k*th fold. The cross-validation estimate of the expected test error is:

$$CV\left(\hat{f}\right) = \frac{1}{N}\sum_{i=1}^{N} L\left(y_i, \hat{f}^{-k(i)}\left(x_i\right)\right).$$

Overall, five- or tenfold cross-validation are recommended as a good compromise: see Breiman and Spector (1992), Kohavi (1995) and Guyon et al. (2006).

3.3 Experimental set up of our Stacking proposal

Most of the research on ensembles of classifiers tends to demonstrate that Stacking can perform comparably to the best of the base classifiers as selected by cross-validation, if not better. It would seem natural to expect that the final classifier produced by Stacking is able to achieve performances better in terms of greater accuracy than the best level-0 classifier. Otherwise, the computational onus created by the complexity of the procedure would not be justified.
This has motivated us to investigate the performance of the Stacking technique empirically, also in terms of stability and strength, by proposing a scheme that also modifies some elements among those, firstly, the starting point, which will be represented by a simulation study that will generate datasets with different parameters and different contamination levels. We, therefore, focus on an aspect that has been covered much less by the studies and that in our view, however deserves special attention: the choice of the initial dataset. This is connected to the assumption that small changes in the dataset can lead to different models and that the presence of outliers might then alter the parameters of the model.

Our thesis proposes, with regard to the "traditional" approach, the following objectives:

1) Evaluation of the base-level and meta-level classifiers in terms of their accuracy when there are modifications in the input data set

2) Evaluation of the effects caused by the presence of anomalous values in the data set on the performances of the base-level and meta-level classifiers and their comparison

3) Evaluation of the results of the simulation studies carried out to establish whether, and to what extent, the combination of classifiers makes it possible to improve performances compared to the use of a single classifier.

In order to achieve these objectives we built a Stacking scheme that includes some differences and characteristic implementations compared to what is proposed by the literature, and these will be specifically explained below for each element of the process.

3.3.1 Software

Among the main differences from the predominant literature we find the choice to build the Stacking procedure entirely in a Matlab environment. [MATLAB 7.12.0 (R2011a) and (R2011b)]. This allowed us to create a flexible and powerful work tool, although it is sometimes a little heavy in computational terms.
The whole Stacking scheme has therefore been implemented in Matlab and built in the form of an interlinked process which is begun by the generation of the level-0 data in the case of the simulation study, and which includes a complete homogenisation of the procedures relative to each of its phases in order to guarantee uniformity and therefore comparability of the outputs returned at every step.
Since the Stacking algorithm, although included in other software (such as the open source software WEKA), did not exist in Matlab, a code for creating a complex structure was completely implemented, which made possible:

- The organisation of a complex and extended experimental plan in order to carry out a wide-ranging simulation study;
- The building of a set of base classifiers and the appropriate procedure of cross-validation for carrying out the fit, the assessment and the generation of predictions for the formation of the input dataset for the meta-learner
- The fit and the assessment of the meta-classifier, by means of an appropriate cross validation procedure, and thus the procedures for the prediction combination and the homogeneous processing of the results with regard to the characteristics of each method
- The creation of suitable plots

It was necessary to carry out a process of homogenisation for each step of the procedures for all the classifiers, as we pointed out, which were chosen voluntarily with different characteristics, in order to obtain the same output that is indispensable for making the structure function and for the assessment and comparability of the results.

Of course this has also led to homogeneity for the base classifiers in the procedures for the entire construction process for each one, to the generation of the predicted class labels and the relative computation of the prediction error (in this case cross validation error or extra-sample error which represents the fraction of the misclassified observations of the test total computed by the difference between the predicted class label and the true class label relative to the test set), which meant, earlier in the process, the implementation of the stratified k-fold cross-validation procedure, constructed in the same way for all the classifiers, including those for which this was not planned, which allowed us to achieve the same data partition in training sets and test sets. Similarly, since we follow the approach of using probability distributions generated by base classifiers as metadata (because we believe that this is better than the predictions), but not all the selected base classifiers return class probabilities as output, we implemented the Matlab procedure for each classifier, in order to generate posterior probabilities (or to make transformation from predictions to probabilities), which

are indispensable for creating meta-classifier input datasets, especially for those models which are not predicted by default. Thus the classifiers were not chosen on the basis of return output, but on the basis of their heterogeneity, which is a contribution to their knowledge of the phenomenon.

With regard to the decision rules and the typical characteristics of each algorithm, the following have been made available for each classifier, and implemented if not already present in Matlab :

- Posterior probabilities of training
- Training error rate
- Predicted class labels
- Posterior probabilities of testing obtained through ten-fold cross-validation
- Mechanisms for the partition and iteration of the dataset for cross validation
- Cross validation error rate

3.3.2 Simulation study for 0-level input data

The proposals from the literature reported so far have always used well-known datasets from the UCI learning repository for the building of Stacking schemes (Blake and Merz 1998)) .
With reference to what has been pointed out several times, that is to say that even small changes to the training set may lead to different models and, taking into account that little has been is known about the relationship between training dataset parameter settings and performances of base classifiers and meta-classifiers, we chose not to use well-known datasets, contained in databanks and used extensively for dealing with problems of supervised classification.
The part that deals with the data is, in our opinion, an important moment in our study, and taking this into consideration, we carried out a wide-ranging simulation study which led to the generation of datasets with different characteristics as follows:

$N = size\ of\ the\ population$

$k = number\ of\ classes$

$nj = size\ of\ the\ class \qquad (j = 1, ..., k)$

$v = number\ of\ features$

$\delta = degree\ of\ separation\ among\ the\ subpopulations$

$prior = prior\ probability\ to\ belong\ to\ a\ class$

$Niter = number\ of\ iterations$

For each dataset we generated two groups of *nj* observations; the first group consists of a *nxv* matrix generated from a standard multivariate normal population with a mean equal to μ for all

variables and a covariance matrix Σ. The Second Group was also generated from a multivariate normal population, but with a mean equal to $(\mu + \delta)$.

We therefore imposed a different degree of separation δ between the groups.
Furthermore, taking into consideration the effects that atypical observations might have on the model too, we decided to build a well-organised and complex experimental contamination design that would allow us to explore this particular aspect of the problem, which has been paid very little attention by the literature, and draw some conclusions.

Contamination design

Typology: shift contamination
Level: proportion of contaminated data:
The contamination can be carried out on:

- only one class
- on both

Intensity: value of the constant to be added to the original data.

Both for datasets with and without contamination, each simulation in the Stacking process is repeated in a series of 100 trials. In each trial a dataset is generated by following the characteristics indicated in the section above.

3.3.3 Base Classifiers

The proposed Stacking scheme includes the use and subsequent combination of a set of 13 base-level classification methods generated by the 8 learning algorithms, as indicated in the subdivision of section 3.2.1, whose output will be used as input data for the meta-level algorithms.
The set is therefore composed of classifiers that have been generated by following several criteria:

- by applying different learning algorithms to a single data set
- by applying a single learning algorithm with different parameter settings to a single data set
- in the case of multiple classifiers by applying a single learning algorithm to the different variants of a dataset (bagging, boosting)

As we have already indicated, very few modifications were made intentionally to their default parameter settings and the exceptions will be included in the description of the single classifiers.
In the following part of this Section we illustrate the main parameter settings of the base classifiers used if they are different from the default ones.

- **Linear Discriminant Analysis and Quadratic Discriminant Analysis**

No significant modifications were made to their default parameter settings, but of course this is without considering what has been indicated for the organisation of the cross validation procedure (which holds for all the classifiers, excluding a version of Classification Tree and of Adaboost, as we

will see later), which has made it possible to return the output predicted class labels and posterior probabilities homogeneously, and also to compute cross validation error.

- **Naive Bayes.**

Since the algorithm provides support for Gaussian and Kernel distribution, both were used in the experimental phase. In fact, it seems appropriate to use the Gaussian distribution for features that have normal distributions in each class, since for each dataset we had generated two groups of *nj* observations consisting respectively of a *nxv* matrix generated from a standard multivariate normal population. However, in the algorithm's training phase we also used a Kernel distribution that is appropriate for features that have a continuous distribution. Since this requires more computing time and more memory than the normal distribution and since in our case the results did not seem significantly better, we preferred to use the normal distribution of the proposed scheme.

- **Classification Tree**

These were used in the two versions:

- with pruning, which computes the optimal sequence of pruned subtrees (TRE)
- without pruning, which returns the decision tree that is the full one (TRE1)

- **Bagged Tree**

To estimate the prediction error of the bagged ensemble, instead of computing predictions for each tree on its out-of-bag observations, we average these predictions over the entire ensemble for each observation and then compare the predicted out-of-bag class value with the true class at this observation (as by default), and we built the cross validation procedure on the entire dataset.
We created an ensemble of 30 bagged decision trees

- **Adaboost**

We use two ensemble algorithms:

- First, (ADAm) based on AdaBoostM1 (Freund and Schapire, 1996). The base classifier returns a discrete class label. Weak learner = tree.
  Number of ensemble learning cycles = 30

- For the second (ADA), we created a personalized function that extends the Matlab function "adaboost", with the implementation of the cross validation procedure, and, for calculating posterior probabilities extends the calibration of the output of AdaBoost. MH proposed by Busa-Fekete et al. (2011) for multi-class problems.
  Number of ensemble learning cycles=30

- **Support Vector Machine (LIBSVM)**

The Support Vector Machines (SVM), developed in the 1990s (Boser et al., 1992; Cortes and Vapnik,1995) are held to be among the most effective methods of supervised learning. They were implemented in the scheme proposed through LIBSVM by Chang and Lin (2011), one of the most widely used SVM software programs. Four different implementations of the algorithm were created and for all of them the transformation of the design matrix was implemented in a sparse matrix, the procedure of common cross validation as for the other classifiers, together with the computing of posterior probabilities for extending SVM to give probability estimates (instead of only class labels as default). The Kernel function was chosen as a reference: RBF (Gaussian) kernel: $K\left(x,x\right)=e^{-\gamma\|x-y\|^2}$, $\gamma>0$

The specific implementations for each version of the algorithm are summarised as follows:

- SVM
- Scaled SVM
  The authors recommend linear scaling. We have chosen to scale each attribute to the range [0,1]
- SVMbest
  A procedure of cross validation was implemented in order to choose the best parameters (C,γ) for an RBF kernel. Various pairs of (C,γ) values are tried and the one with the best cross validation accuracy is chosen. We recommend trying exponentially growing sequences of parameters C and γ to select good ones (e.g. $C=2^{-5},2^{-3},...,2^{15}$ $\gamma=2^{-15},2^{-13},...,2^{3}$). Although the grid search in cross validation is recommended, it means a great deal of computational time, at least with the values suggested.
- scaled SVMbest
  There are the implementations for SVM best and for scaled SVM.

3.3.4 Meta–classifiers

Since we believe that both MLR and ridge regression are valid algorithms for combining the outputs of base classifiers, we decided to use both of them with a mechanism that establishes in a mutually exclusive way the application of linear regression in the hypothesis in which there is no multicollinearity for input matrix X, and of ridge regression when the dimensionality of the meta-feature space L (L=number of base classifiers) is greater than the effective rank of the input matrix.

- **MLR** (Multi-Response Linear Regression)
- **Ridge Regression**

The use of cross-validation on meta-data has also been envisaged for meta-classifiers, to build and then evaluate the meta-classifier, reduces the risk of overfitting and enables us to consider the

estimate of the prediction error given by such a process as a generalization error of the Stacking scheme.
For this part of the thesis, therefore, we made extensive use of cross validation, since we use it in order to:

- build the classifiers (base and meta-level) from the training data
- estimate the prediction error of the base classifiers and the final model
- estimate the unknown tuning parameters (particularly for Ridge Regression and Support Vector Machine).

However, as we have mentioned before, we preferred not to proceed with an extreme tuning of the parameters with regard to the objective of investigating whether the Stacking with the combination of different methods is able to improve the performances of the classifiers, mitigating any bad performances, especially those of the "weakest" ones.

3.3.5 Evaluating and Comparing classification methods

The generalization errors of the base classifiers for a given input dataset and of a meta-classifier (for a input dataset generated from partial class probability distributions from each base classifier) are estimate by averaging the result of 100 runs of ten-fold stratified cross validation . Cross validation is repeated 100 times using different random seeds of the data resulting in 100 different sets of folds. The same folds are used in all experiments to built all the base classifiers and to estimate their true errors.
It should always be remembered that even though we put together values relative to Stacking and base classifier errors in the tables and in the different plots, they are constructed using different typologies of input data. A comparison can be made among Stacking schemes and for the single Stacking scheme just to establish whether or not Stacking is the better, worse or at least equal to the best base classifier
Because of the variability and fluctuation of the cross validation error, the average does not seem to be sufficient and in addition other measures are calculated on the distribution of the cross validation errors for each classifier and averaged over the total of the iterations carried out.

- **Position indices**
  - Median
  - Percentage of best positioning

- **Indices of variability**
  - Standard deviation
  - Median of deviations from the median
  - Interquartile difference
  - Range

3.4 Discussion

The state of the art in the research of the Stacking framework was a valid starting point for this thesis. In this chapter we have illustrated the main characteristics of the proposed Stacking scheme, starting from the choice of implementing the whole Stacking scheme in the Matlab environment and built in the form of an interlinked process which is begun by the generation of the level-0 data in the case of the simulation study which led to the generation of datasets with different characteristics and, furthermore, taking into consideration the effects that atypical observations might have on the model too, to build a well-organised and complex experimental contamination design that would allow us to explore this particular aspect of the problem, to which very little attention has been paid by the literature, and to draw some conclusions. Of course, the implementation of the entire Stacking scheme required a complete homogenisation of the procedures relative to each of its phases in order to guarantee uniformity and therefore a comparability of the outputs returned at every step. Furthermore, the creation of a double procedure of cross-validation both for base and meta-classifiers (which are represented by MLR and Ridge Regression, the use of which is regulated by a mutually exclusive insertion mechanism where conditions of collinearity occur), made it possible to build and evaluate classifiers at a double level, thus reducing the risk of overfitting too. Together with cross validation error, other measures have been included, calculated on the error distribution of cross validation for each classifier and for Stacking. In the following chapter ample space will be given to the results of the application of the proposed scheme to the datasets generated by the experimental design, contaminated and non-contaminated, and on real datasets to empirically verify their functioning. With a view to a further improvement of the entire proposed process, which is at the experimental stage, the research activity will be directed towards optimising performances and guaranteeing the reliability of the predictions for single classifiers by modifying the setting of the parameters used in this phase, and more generally by including:

- Extension of the experimental design both to verify further the results achieved and to insert other elements into the design (different processes of data generation, increasing the number of classes, different *prior* values)
- Introduction of more adequate measurements (at some point combining the use of more than one index) which are able to capture accuracy in the best possible way in terms of estimating the prediction error returned by the classifiers
- Possible introduction of a weighting system into the method of meta-classification should we intend to combine several classifiers with very different performances in terms of accuracy.
- Extension of the methods proposed by the literature (Varma et al., 2006 ; Tibshirani et al., 2008) for the estimation and reduction of potential *bias* in cross validation error for the problem that is the object of the thesis

CHAPTER 4

Experimental results

4.1 Introduction

In this section there is a summary of the main results obtained from the application of the proposed Stacking scheme to datasets generated by means of the experimental design and also real data. In particular, in section 4.2 the results relative to the non-contaminated data are illustrated, in order to investigate the effects on the performance of the base classifiers and of Stacking in the presence of input datasets with different characteristics. In section 4.3 the application was carried out on simulated and contaminated data to investigate whether the presence of outliers can affect the performances of the base classifiers and of Stacking. In Section 4.4 are illustrated the results obtained using three different Stacking variants with different base-level classifier subsets built on different datasets In Sections 4.5 and 4.6 there are applications to real data.

4.2 Simulated data

Based on the characteristics indicated in the previous chapter, datasets have been generated with the following characteristics, which make them different in terms of complexity and degree of separation among the groups:

N = 120; 200
n_1 = 60; 100
n_2 = 60; 100
v = 3; 5 ; 7 ; 10
δ = 0.5; 1; 1.5; 2; 2.5; 3

The experimental design for a fixed scheme with 13 base-level classifiers is used, to which we always refer for completeness. Subsequently, some examples of level-0 data will be reported for the base-level algorithms. Ten-fold stratified cross-validation was used on each dataset to build single methods and estimate the prediction error of all the base-level classifiers. Every trial process was repeated 100 times and an average of the results was calculated in order to find the error of average cross validation for each classifier, relative to each experiment. The posterior probabilities of each classifier derived from the process of cross-validation form the meta-dataset for the meta-classifier.
Ten-fold stratified cross-validation is also used, repeated 100 times, for the meta classifiers, which are Linear regression and Ridge regression (mutually exclusive when hypotheses of multicollinearity recur).
We are interested in investigating empirically the performances of the single classifiers and of Stacking for datasets with different characteristics and above all if it is convenient to use

Stacking in terms of improving performances instead of a single classifier, bearing in mind the necessary increase in computation.
It should always be remembered that even though we put together values relative to Stacking and base classifier errors in the tables and in the different plots, they are constructed using different typologies of input data. A comparison can be made among Stacking schemes and for the single Stacking scheme just to establish whether or not Stacking is better, worse or at least equal to the best base classifier.
An analysis of the results obtained by applying the Stacking scheme to the set of the datasets generated by means of the experimental design does not lead us to believe that the prediction error of the Stacking scheme is to be considered lower than any other classifier or that, therefore, the Stacking scheme is preferable in terms of performances to the use of the best single classifier.
It always achieves good performances and is to be considered among the best, but it does not seem to be preferable for this type of application.
As we can see in an example summarised in table 4.2.1, which shows the results of the application of the Stacking scheme to the dataset obtained from the experimental design characterised by $n_1 = n_2 = 60$, $v = 10$, $\delta = 0.5$. It should be noted that the best positioning is always in agreement with the lowest error.

	Cross Validation Error	Median Cross Validation Error	Std. Deviation Cross Validation Error	Difference interquartile Cross Validation Error	Range Cross Validation Error	MAD Cross Validation Error	% Best positioning
Classifier							
LDA	0,2334	0,2300	0,0315	0,0450	0,1450	0,0200	12
QDA	0,2703	0,2700	0,0398	0,0550	0,2000	0,0300	0
TRE	0,3565	0,3550	0,0418	0,0500	0,2350	0,0250	0
TRE1	0,3580	0,3600	0,0386	0,0450	0,2050	0,0250	0
BAG	0,2799	0,2800	0,0347	0,0475	0,1750	0,0250	2
ADA	0,2767	0,2800	0,0357	0,0475	0,1900	0,0250	1
ADAm	0,2764	0,2750	0,0340	0,0475	0,1500	0,0250	0
NBA	0,2353	0,2350	0,0287	0,0400	0,1200	0,0200	12
SVM	0,2443	0,2400	0,0329	0,0425	0,1500	0,0225	7
SVMscaled	0,2239	0,2200	0,0315	0,0450	0,1400	0,0200	33
SVMb	0,2443	0,2400	0,0329	0,0425	0,1500	0,0225	0
SVMbscaled	0,3497	0,2650	0,1542	0,3025	0,4650	0,0650	9
GLM	0,2313	0,2300	0,0300	0,0450	0,1400	0,0200	8
STA	0,2314	0,2350	0,0361	0,0450	0,1850	0,0225	16

Table 4.2.1. Simulated Data.Measures of the performances of the classifiers and of the Stacking scheme. Average figures over 100 iterations. N=200, v=10,δ=0.5

We can appreciate a certain variability in the boxplots of error distribution in Figure 4.2.1.

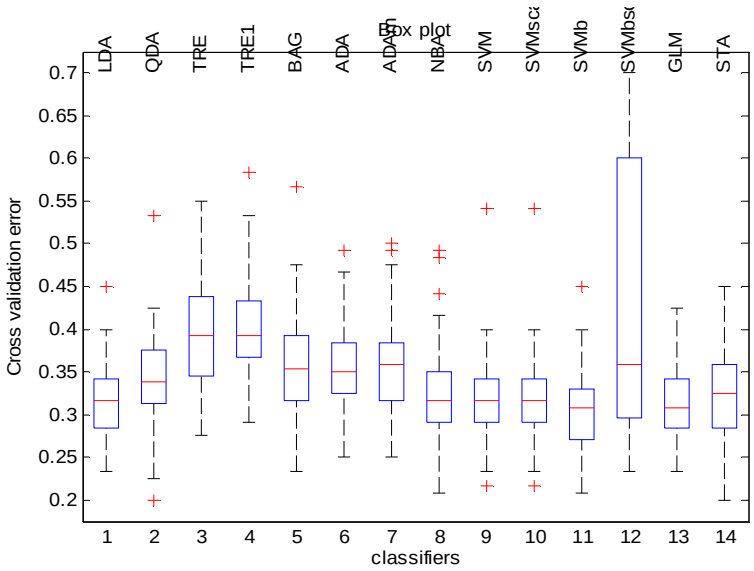

Figure 4.2.1. Simulated Data. Boxplots of error distribution on 13 base classifiers and Stacking scheme. N=200, v=10,δ=0.5

On a more general level, Table 4.2.2 illustrates the results relative to the average cross validation error for each base classifier and for Stacking for input datasets with varying degrees of complexity and with different degrees of separation between the two groups.

	δ=0.5				δ=3			
	n. variables				n. variables			
Classifier	3	5	7	10	3	5	7	10
LDA	0,3432	0,3145	0,2738	0,2396	0,0053	0,0003	0,0001	0,0000
QDA	0,3539	0,3392	0,3123	0,2948	0,0058	0,0003	0,0002	0,0000
TRE	0,4174	0,3958	0,3807	0,3602	0,0392	0,0409	0,0428	0,0418
TRE1	0,4243	0,3980	0,3761	0,3603	0,0381	0,0397	0,0425	0,0407
BAG	0,3921	0,3586	0,3258	0,2840	0,0197	0,0101	0,0035	0,0023
ADA	0,3785	0,3553	0,3213	0,2809	0,0201	0,0193	0,0199	0,0223
ADAm	0,3795	0,3557	0,3177	0,2863	0,0206	0,0048	0,0040	0,0072
NBA	0,3509	0,3236	0,2809	0,2476	0,0050	0,0003	0,0000	0,0000
SVM	0,3436	0,3188	0,2856	0,2431	0,0058	0,0005	0,0001	0,0000
SVMscaled	0,3347	0,3032	0,2657	0,2274	0,0044	0,0002	0,0000	0,0000
SVMb	0,3436	0,3188	0,2856	0,2431	0,0058	0,0005	0,0001	0,0000
SVMbscaled	0,4033	0,4252	0,4069	0,4119	0,0428	0,0052	0,0076	0,0000
GLM	0,3427	0,3129	0,2757	0,2428	0,0092	0,0009	0,0002	0,0000
STA	0,3545	0,3237	0,2846	0,2467	0,0073	0,0012	0,0001	0,0001

Table 4.2.2. Simulated Data. Cross validation error for each base classifier and for Stacking for input datasets with varying degrees of complexity and with different degrees of separation between the two groups. N=120

The analysis was carried out on the entire set of the generated datasets, but for the sake of brevity we will only report the most important results. The behaviour of scaled SVM (which represents one of the implementations adopted for the Support Vector Machine described in Chapter 3) is particularly interesting , as it is has the lowest error among the classifiers, also in comparison with Stacking. Instead SVMbest does not achieve such moderate error levels and this pushes us to improve the cross validation procedure used for tuning the parameters. Thus, by

using a set of classifiers containing scaled SVM, Stacking scheme could prove itself not to be competitive, as we can also see in Table 4.2.3 which summarises the performances of scaled SVM and the Stacking scheme on all the datasets generated by this part of the experimental design.

	Stacking scheme					SVMscaled			
	n. variables					n. variables			
δ	3	5	7	10	δ	3	5	7	10
0,5	0,3545	0,3237	0,2846	0,2467	0,5	0,3347	0,3032	0,2657	0,2274
1	0,2095	0,1515	0,1080	0,0691	1	0,1945	0,1388	0,0971	0,0608
1,5	0,1074	0,0576	0,0316	0,0153	1,5	0,0972	0,0470	0,0245	0,0093
2	0,0507	0,0187	0,0071	0,0019	2	0,0402	0,0124	0,0042	0,0006
2,5	0,0201	0,0035	0,0012	0,0001	2,5	0,0139	0,0018	0,0006	0,0000
3	0,0073	0,0012	0,0001	0,0001	3	0,0044	0,0002	0,0000	0,0000

Table 4.2.3.Simulated Data. Comparison between the cross validation error given by scaled SVM and the Stacking scheme on the datasets generated by the experimental design. N=120, different level of degree of separation and number of variables

On going back to table 4.2.2 we notice that apart from the above-mentioned very good behaviour of scaled SVM, generally speaking the parametric classifiers reach some of the best levels of accuracy, while "weak" classifiers, such as Classification Tree and the ensemble methods such as Bagged Tree and Adaboost achieve rather disappointing performances.

	δ=0.5				δ=3			
	n. variables				n. variables			
Classifier	3	5	7	10	3	5	7	10
LDA	16	14	17	15	75	98	99	100
QDA	4	4	3	1	9	1	0	0
TRE	2	0	0	0	1	0	0	0
TRE1	1	0	0	0	0	0	0	0
BAG	1	5	3	1	2	0	0	0
ADA	7	1	1	3	3	0	1	0
ADAm	3	7	6	2	0	0	0	0
NBA	12	9	8	8	4	0	0	0
SVM	9	8	4	15	0	0	0	0
SVMscaled	22	27	27	27	3	1	0	0
SVMb	0	0	0	0	0	0	0	0
SVMbscaled	5	5	7	7	1	0	0	0
GLM	6	6	13	11	1	0	0	0
STA	12	14	11	10	1	0	0	0

Figure 4.2.4. Simulated Data. Best positioning over 100 iterations for each classifier and Stacking scheme for different level of degree of separation and number of variables. N=120

As far as the analysis of results relative to the measure that counts the number of times an algorithm performs better than the others (over 100 iterations) is concerned, Figure 4.2.4 shows that for a low level of degree of separation, scaled SVM achieves the best relative positioning compared to the other classifiers, while with a higher degree of separation LDA achieves by far the best positioning.

This result is also confirmed for N=200. For datasets with intermediate degrees of separation (≥1.5) the best positioning is always achieved by LDA, above all if associated (in the lower values) with a higher complexity due to the number of variables.
By comparing Table 4.2.2 and Table 4.2.4 (and more generally the whole set of results) we can observe that there is not always a correspondence for the classifiers between the best positioning achieved and the lowest value of cross validation error achieved and also because, since we are dealing with an average of 100 iterations, the variability of a classifier is significant in terms of the error returned which is often very high. It may therefore happen that classifiers with a higher variability can achieve better positioning.
This circumstance should lead to us to continue looking for more adequate measurements (at some point combining the use of more than one index) which are able to capture accuracy in the best possible way in terms of estimating the prediction error returned by the single classifiers also in order to improve the comparison with the use of a more complex scheme like Stacking.

4.3 Simulated contaminated data

The second application of the Stacking procedure to simulated data was carried out on datasets generated with the same characteristics as the ones used in the previous section but they have undergone a contamination design which included a shift contamination with different levels of contaminated data, carried out on only one class or both and with different values of the constant to be added to the original data. Let's summarize the main characteristics of the contamination design:

Level (proportion) of contaminated data: 5%; 10%; 30%
Number of classes : one; both
Value of the constant to be added: +2; +4

They were generated by means of the contamination design and about circa 480 datasets were analysed.
In this section we are interested in investigating empirically the performances of the single classifiers and of Stacking for datasets upon which a contamination has been carried out and above all in seeing if it is convenient to use Stacking in terms of improving performances instead of a single classifier, bearing in mind the necessary increase in computation.
In the case of contaminated data Stacking improves its performances noticeably compared to what we have observed for non-contaminated data, in some cases also in comparison with scaled SMV, and generally appears to be very competitive, above all when the contaminations are more substantial. In a set of classifiers in which there was no SVM it would be the best for each of the analysed datasets. While scaled SVMbest achieves cross validation error values that are always very substantial and therefore are definitely not to be inserted with the current parameter setting (selected by means of cross validation) in a basic set of classifiers. It will be advisable to improve the tuning of the parameters if the decision is taken to use it. It was inserted in order to make the analysis complete, but because of its very bad performances, it will not be taken into account in the comparison with the other classifiers. We will distinguish the hypotheses in which the contamination on the total of the observations is moderate (10%) from the hypotheses in which it is stronger (30%) and is carried out on one or both the classes. Table 4.3.1 illustrates the results relative to the average cross validation error for each base classifier and for Stacking for different levels of contamination with different degrees of separation

between the groups and for different number of variables of the input datasets. The contamination is only intended for one class: cont=+ 4. N=120.

	10%				30%			
	δ=0.5		δ=3		δ=0.5		δ=3	
	n. variables				n. variables			
Classifier	3	10	3	10	3	10	3	10
LDA	0,5092	0,5007	0,1211	0,1220	0,5083	0,4667	0,4667	0,4625
QDA	0,3789	0,3377	0,1210	0,1098	0,5250	0,5667	0,2240	0,1714
TRE	0,3853	0,3502	0,1456	0,1397	0,3750	0,3167	0,2416	0,2161
TRE1	0,3903	0,3502	0,1426	0,1446	0,3833	0,3500	0,2426	0,2186
BAG	0,3631	0,2708	0,1131	0,0698	0,2583	0,2583	0,2098	0,1045
ADA	0,3845	0,3618	0,1450	0,1414	0,3083	0,2750	0,2513	0,2073
ADAm	0,3531	0,3066	0,1337	0,1051	0,2583	0,2750	0,2235	0,1672
NBA	0,3945	0,3928	0,1243	0,1009	0,5000	0,5000	0,2715	0,2334
SVM	0,3073	0,2201	0,0858	0,0363	0,2250	0,1250	0,1590	0,0597
SVMscaled	0,3022	0,2067	0,0788	0,0333	0,2417	0,1167	0,1541	0,0510
SVMb	0,3073	0,2201	0,0858	0,0363	0,2250	0,1250	0,1590	0,0597
bscaled	0,4873	0,4733	0,1352	0,0930	0,4917	0,4750	0,4153	0,3868
GLM	0,5072	0,4994	0,1362	0,1597	0,5083	0,4667	0,4728	0,4646
STA	0,3248	0,2145	0,0876	0,0362	0,2417	0,1167	0,1681	0,0552

Table 4.3.1. Simulated contaminated data. Cross validation error for 13 base classifier and for Stacking for different levels of contamination, different degrees of separation between the groups and for different number of variables of the input datasets. The contamination is only intended for one class: N=120.

When only one class is contaminated with a proportion of 10%, the behaviour of Stacking is very good and is only exceeded by three implementations of SVM. On the contrary, the effect of the contamination is quite substantial for Linear Discriminant Analysis and Logistic Regression, as well as for scaled SVMbest, which we have already said will no longer be inserted into the set of classifiers, as we prefer SVM best or scaled SVM. By increasing the degree of separation between the groups, we find quite uniform behaviour for the datasets with a moderate number of variables, with the exception of SVM and Stacking, while the application to a dataset with a higher number of variables greatly improves the performances of Naive Bayes, of the three best implementations of SVM and of Stacking. We will see that by contaminating only one class with a level of 30% there will be a very clear effect on the four parametric base classifiers present in the set and in particular on the behaviour of NBA which, in a contamination hypothesis of 30% of the data in the presence of a low degree of separation between the groups will return a fixed error of 0.50 for every iteration. Stacking achieves an error level equal to that of the best classifier, which is scaled SVM. In the presence of a higher degree of separation (δ=3), the worst performance is given dramatically only by LDA and Logistic Regression.

Moving on to the hypothesis of contaminating two classes, in Table 4.3.2 we can see that with a moderate contamination and a low degree of separation between the classes, the four parametric classifiers in any case achieve a high level of error, although they are not among the worst.

	10%				30%			
	δ=0.5		δ=3		δ=0.5		δ=3	
	n. variables				n. variables			
Classifier	3	10	3	10	3	10	3	10
LDA	0,4254	0,4583	0,0883	0,1094	0,4667	0,4583	0,3288	0,3033
QDA	0,4633	0,4873	0,0927	0,1843	0,4667	0,5750	0,3623	0,3552
TRE	0,4306	0,3973	0,1054	0,1195	0,4750	0,3833	0,1607	0,1610
TRE1	0,4303	0,3881	0,1036	0,1188	0,4417	0,3917	0,1564	0,1558
BAG	0,4002	0,3053	0,0736	0,0525	0,4250	0,3000	0,1352	0,0845
ADA	0,4109	0,3723	0,0904	0,0904	0,3917	0,3583	0,1605	0,1584
ADAm	0,3951	0,3547	0,0903	0,0748	0,4083	0,4000	0,1554	0,1368
NBA	0,4309	0,4373	0,0758	0,0509	0,5000	0,5000	0,3921	0,4523
SVM	0,3610	0,2357	0,0543	0,0268	0,3417	0,2250	0,0966	0,0383
SVMscaled	0,3479	0,2379	0,0489	0,0234	0,3667	0,2333	0,0927	0,0331
SVMb	0,3610	0,2357	0,0543	0,0268	0,3417	0,2250	0,0966	0,0383
bscaled	0,4494	0,3653	0,2683	0,0712	0,5667	0,3750	0,2195	0,2116
GLM	0,4204	0,4524	0,0703	0,0873	0,4500	0,4500	0,2563	0,2689
STA	0,3532	0,2377	0,0580	0,0245	0,3667	0,2167	0,1020	0,0376

Table 4.3.2. Simulated contaminated data. Cross validation error for 13 base classifier and for Stacking for different levels of contamination, different degrees of separation between the groups and for different number of variables of the input datasets. The contamination is intended for two class: N=120. Cont=+4.

With a higher level of contamination for the two classes and a low degree of separation between the groups, they are always among the worst, and generally speaking none of the classifiers gives a good performance. The behaviour of Stacking is always very interesting and overall is preferable for (δ=0.5, v=10), thus confirming the analogous result obtained also in the case of the contamination only of a single class in Table 4.3.1. In the case of a higher degree of separation (δ=3), the worst performance is achieved once again by Linear Discriminant Analysis and Logistic Regression.
Generally speaking, based on the outcome of the experiments carried out, it seems to be that the contamination contained by a single class causes a deterioration in the performances of LDA and Logistic Regression, while with a higher level of contamination this only happens with a high degree of separation between the groups. Where there is a low degree of separation, the worst performance will be achieved for all four of the parametric classifiers. The behaviour of Naive Bayes is particularly interesting, since at a high level of contamination, carried out both on one class and on both using the hypothesis of a low degree of separation, it always obtains a value equal to 0.50 which remains the same for each iteration as can be seen in Figure 4.3.1, so using it is not very effective.

Table 4.3.3 summarises the results relative to the average cross validation error for each base classifier and for Stacking for different levels of contamination with different degrees of separation between the groups and for different numbers of variables of the input datasets. The contamination is carried out on both the classes N=200

	10%				30%			
	δ=0.5		δ=3		δ=0.5		δ=3	
	n. variables				n. variables			
Classifier	3	10	3	10	3	10	3	10
LDA	0,4171	0,4443	0,0817	0,0861	0,4900	0,4659	0,3330	0,3199
QDA	0,4498	0,4807	0,0849	0,1329	0,4700	0,4792	0,3702	0,3470
TRE	0,4255	0,3592	0,0974	0,0943	0,4200	0,3419	0,1466	0,1357
TRE1	0,4215	0,3633	0,0981	0,0947	0,4450	0,3427	0,1483	0,1381
BAG	0,3990	0,2860	0,0684	0,0461	0,4750	0,2520	0,1230	0,0669
ADA	0,4017	0,3437	0,0884	0,0763	0,4000	0,3193	0,1451	0,1221
ADAm	0,3955	0,3307	0,0882	0,0656	0,4200	0,3226	0,1394	0,1148
NBA	0,4236	0,4400	0,0734	0,0505	0,5000	0,5000	0,3934	0,4556
SVM	0,3513	0,2299	0,0523	0,0234	0,3250	0,1992	0,0908	0,0330
SVMscaled	0,3489	0,2240	0,0495	0,0211	0,3400	0,1888	0,0891	0,0294
SVMb	0,3513	0,2299	0,0523	0,0234	0,3250	0,1992	0,0908	0,0330
bscaled	0,4375	0,3922	0,1483	0,0501	0,5650	0,5088	0,1987	0,1706
GLM	0,4130	0,4347	0,0649	0,0666	0,4800	0,4636	0,2562	0,2609
STA	0,3497	0,2220	0,0537	0,0231	0,3100	0,1850	0,0959	0,0316

Table 4.3.3. Simulated contaminated data. Cross validation error for 13 base classifier and for Stacking for different levels of contamination, different degrees of separation between the groups and for different number of variables of the input datasets. The contamination is intended for two class. N=200. Cont=+4:

In comparison with Table 4.2.5 with reference to non-contaminated data, we can observe that Stacking does not seem to be affected by contamination in the same way as scaled SVM does. In fact, Stacking is the best classifier in the hypothesis with a low degree of separation, both in the case of contamination at 10% (for v=10) and at 30% . Instead for all the other classifiers the effect is quite considerable, as can be seen in Figure 4.3.1 which can be compared with Figure 4.2.1 relative to the dataset generated by the experimental design but not contaminated.

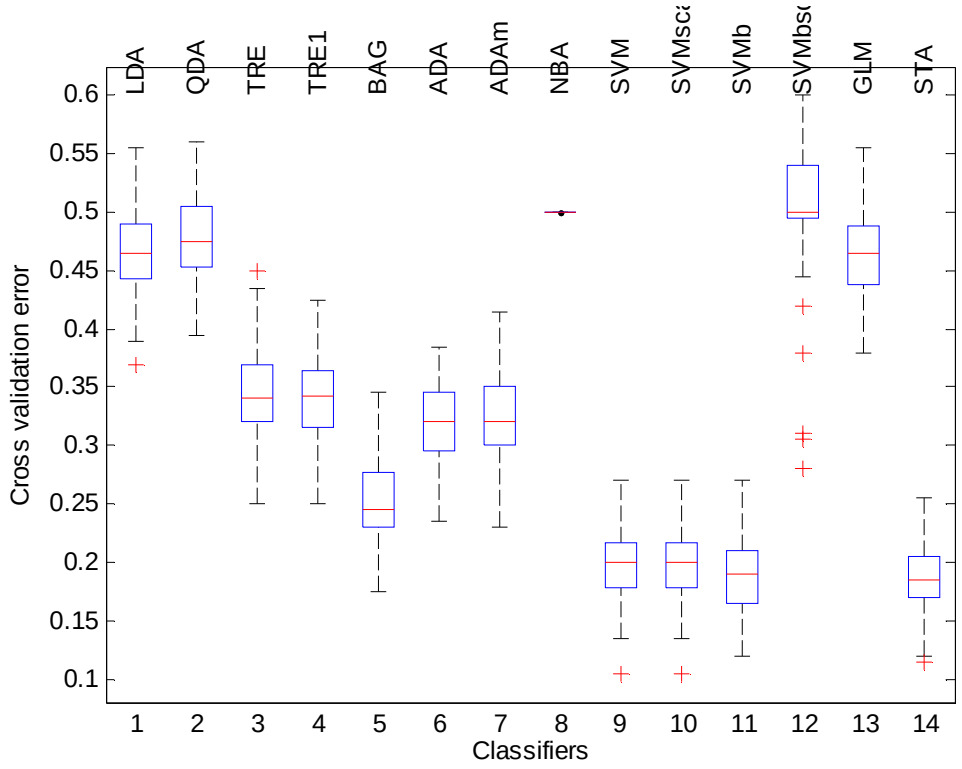

Figure 4.3.1. Simulated contaminated data. Boxplots of error distribution on 13 base classifiers and Stacking scheme. N=200, v=10,δ=0.5, cont=+4, level=30%.

The differences between the performances of the parametric classifiers in the case of contamination of one or both the classes will be re-examined in Chapter 6 with the application of the Stacking scheme in a Forward Search context on simulated and contaminated datasets.

4.4 **Are the number and type of base classifiers important?**

An ensemble of classifiers consists of a set of different classification algorithms and a function to combine their individual outputs in order to improve accuracy.
Any classifier may be used in a Stacking learning scheme and any number of classifiers may be used.
For this thesis we have used a fixed scheme with 13 base-level classifiers.
In this Section, we are interested in investigating whether the performance of Stacking may be modified if we use different base-level classifier subsets on datasets with different characteristics. To this end, we have created three different Stacking variants with different base-level classifier subsets built on different datasets.

The three subsets are differentiated according to the number of classifiers chosen and to their characteristics:

- subset I: composed of 3 classification algorithms : Tree Bagged, AdaboostM1 and SVMbest. The relative Stacking scheme is identified by **"STA3".**
- subset II: composed of 4 classification algorithms: 2 different variants of Classification Tree, Tree Bagged and AdaboostM1. The relative Stacking scheme is identified by "**STA4".**
- subset III: composed of 6 classification algorithms: Linear Discriminant Analysis,Quadratic Discriminant Analysis, Tree Bagged, AdaboostM1, Naïve Bayes Classifier and Logistic Regression. The relative Stacking scheme is identified by **"STA6".**

In the first subset there are classifiers that may be defined as "strong", since both Tree Bagged and Adaboost are themselves members of the ensemble generated by applying a single learning algorithm (a classification tree, which is considered as a "weak" classifier) and the Support Vector Machine (in its version with the selection of the best model through special cross validation) has proved itself to be fairly stable for predictions.

In contrast, in the second subset there are two "weak" learning algorithms and two very stable ones. The reason for this choice is to investigate whether a combination of different strength classifiers can still guarantee satisfactory performances.

For the third subset, which is much more numerous, four parametric classifiers were chosen, two of which are popular but different linear methods for classification tasks (Linear Discriminant Analysis and Logistic Regression). The main difference between them is in the way the linear function is fitted to the training data. Naïve Bayes, which is a simple method, often tends to outperform more sophisticated algorithms when the training set is small.
The experimental plan for the main set of base-level classifiers was used, to which we refer for completeness. Subsequently, some examples of level-0 data will be reported for the base-level algorithms. Ten-fold stratified cross-validation was used for all the base-level classifiers. As in the previous experiments, at every step of cross validation, one part of the available data was used to fit the model, and a different part was used to estimate individual prediction error.
Every trial process was repeated 100 times and an average of the results was calculated in order to find the error of average cross validation for each classifier and relative to each experiment.
The posterior probabilities of each classifier derived from the process of cross-validation form the meta dataset for the meta-classifier.
Ten-fold stratified cross-validation is also used, repeated 100 times, for the meta classifiers, which also in this case are Linear regression and Ridge regression (mutually exclusive when hypotheses of multicollinearity recur).
For completeness, in Appendix to Chapter 4 we show tables with the results of some experiments relating to the performances of the base-level classifier sets on the different dataset input and to the Stacking for different combination schemes. Any slight differences in the performances of some classifiers in the various schemes may be due to different partitions of the cross validation. Stacking turns out to be competitive and better too when compared with other classifiers, especially when there is greater complexity in the base models, or rather when the input datasets are characterised by one dimension larger in terms of observations and variables . The effect of these circumstances is however accentuated when there is a high degree of separation between the two populations.

In Table 4.4.1 and in Figure 4.4.1 some results are summarised relating to the performances of the three different Stacking schemes built with base-level classifier sets of different sizes. The different level-0 datasets in the first column refer to the data input with which the base-level classifiers were built and their prediction errors and probabilities were estimated (via cross-validation) and which form each Stacking scheme. The input data for the meta-classifier, therefore, are always the probabilities generated by base classifiers.

	STA3	**STA4**	**STA6**
Base Classifier Input Dataset			
120_3_05	0.3547	0.3903	**0.3472**
120_10_05	**0.2323**	0.2798	0.2399
120_3_3	**0.0048**	0.0239	0.0066
120_10_3	**0**	0.0059	**0**
200_3_05	0.3498	0.3801	**0.3450**
200_10_05	**0.2268**	0.2732	0.2324
200_3_2	**0.0434**	0.0656	0.0462
200_10_2	**0.0004**	0.0092	0.0016
200_3_3	0.0058	0.0168	**0.0016**
200_10_3	**0**	0.0025	**0**

Table 4.4.1. Cross validation error rate for the different Stacking schemes built

By observing the results, we can say first of all that the STA4 ensemble classification method is the worst for any kind of dataset. The weakness of two of the base-level classifiers was not sufficiently balanced out by the presence of the two ensemble classifiers, which are more stable. In this case, Stacking was unable to fully exploit the predictive capability of the stronger classifiers to compensate for the weakness of the other two.

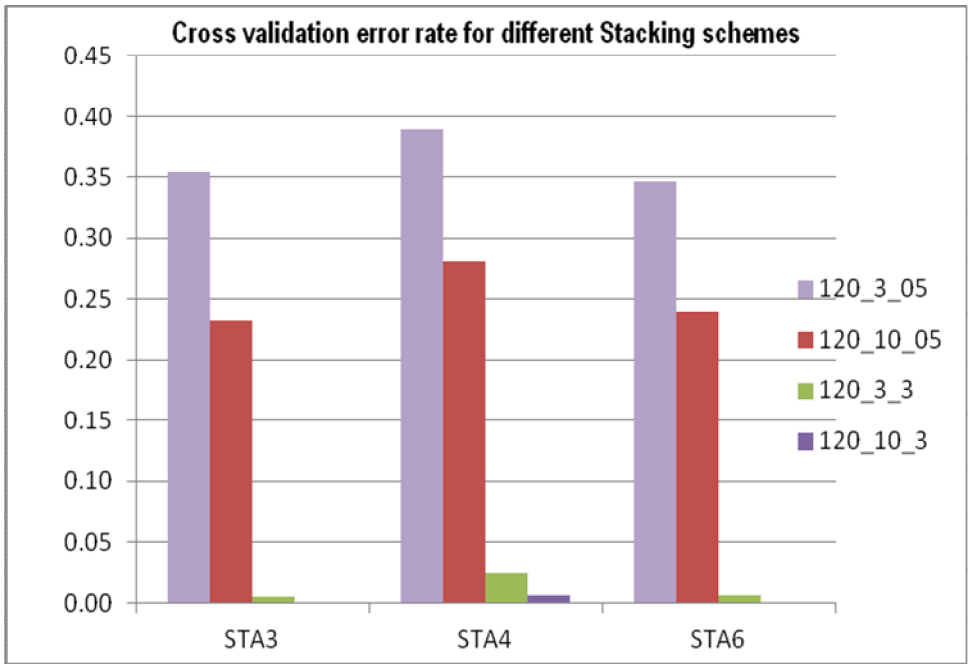

Figure 4.4.1. Cross validation error rate for the different Stacking schemes built and different input dataset

As far as the other two schemes are concerned, there is very little difference between them, although STA3 shows a higher success rate than the more complex STA6.
This might lead us to think that the choice from the start of high performance base-level classifiers could increase the predictive capacities of Stacking.
Furthermore, together with a greater complexity of a combination method built with a greater number of base-level classifiers, such as STA6, there is a better performance, also on level-0 datasets which express less complexity in terms of the number of variables and the degree of overlapping between the two populations in order to signify the weight expressed by the parametric base-level classifiers in their prediction.
However, a comparison shows that STA3 is more successful (although the differences are slight), thus representing the proposal with the best performances, which is also economical in its combination of the presence of classifiers with different characteristics with a lesser complexity of the model.

4.5 Application to real data. Electrodes Data

This part of the thesis is dedicated to the application of the proposed Stacking scheme to real data. The first dataset used is "Electrodes Data" (in Atkinson, Riani and Cerioli (2004)). This is a dataset made up of five variables which represent measurements of electrodes produced by two machines which therefore represent two groups.

$n_1 = 50$
$n_2 = 50$

$y_{1,}y_{2,}y_{5} = diameters$

$y_{3,}y_{4} = lengths$

50 electrodes come from one machine and 50 from the other one.

These data have a particular characteristic: the data have been transformed by subtracting some constants from the variables.
Since this not a dataset generated by means of our experimental design, it is interesting to observe in advance the scatter plot matrix for the variables in Figure 4.5.1.
A certain overlap is noticeable between the two groups, especially for some variables, while in other cases the separation seems to be more substantial, for example for the variable y4. There also seem to be some anomalous values present, but the problems connected to their identification and the possible effects that they may have on the performances of the classifiers will be dealt with in detail in the part that illustrates the Forward Search approach.

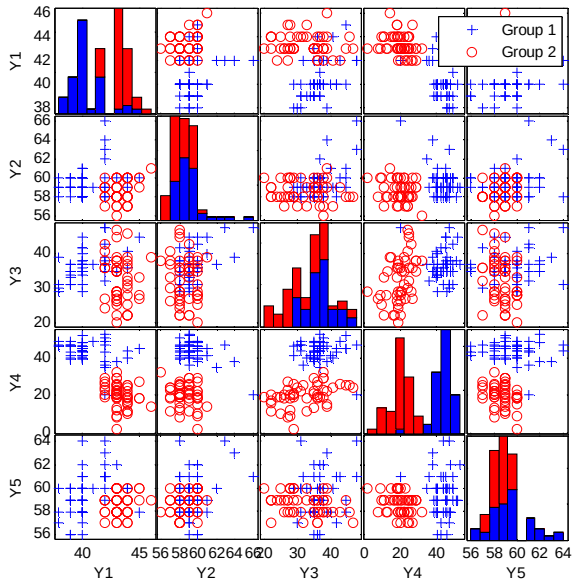

Figure 4.5.1. Electrodes data. Scatterplot matrix with bivariate scatters of the three variables and histograms on the main diagonal. Units in Group 1 are represented by blue crosses.

The main results of the application of the Stacking Scheme to the dataset will be reported below. First of all, in connection with what we covered in the previous section, that is to say the use of different Stacking schemes, we subjected the dataset to the three different Stacking schemes.
In the first scheme we always used a fixed set with 13 base classifiers **STA13,** which represents the scheme of reference that we applied for the simulation study.
On the basis of the experiments performed and the above study carried out for the thesis on the performance of different sets of classifiers, we also decided to use the set with 6 base classifiers **STA6** and the one with three classifiers, **STA3**, with the same composition features reported in the previous section.

We decided to do without STA4, which does not seem to give results that are worse overall compared to the others.
We always used linear least squares regression as meta-classifiers and always ridge regression with the mechanism of mutual exclusivity that has already been described.

4.5.1 Results for STA13

The following table 4.5.1.1 summarizes the results of the measurements of the performances of the various base classifiers and of the Stacking scheme. The figures in the tables are to be understood as an average over 100 iterations of the entire process.

Classifier	Cross Validation Error	Median Cross Validation Error	Std. Deviation Cross Validation Error	Difference interquartile Cross Validation Error	Range Cross Validation Error	MAD Cross Validation Error	% Best positioning
LDA	0,0057	0,0100	0,0052	0,0100	0,0200	0,0000	54
QDA	0,0015	0,0000	0,0036	0,0000	0,0100	0,0000	46
TRE	0,0194	0,0200	0,0024	0,0000	0,0100	0,0000	0
TRE1	0,0192	0,0200	0,0027	0,0000	0,0100	0,0000	0
BAG	0,0142	0,0100	0,0050	0,0100	0,0100	0,0000	0
ADA	0,0317	0,0300	0,0085	0,0000	0,0400	0,0000	0
ADAm	0,0588	0,0600	0,0033	0,0000	0,0100	0,0000	0
NBA	0,0100	0,0100	0,0000	0,0000	0,0000	0,0000	0
SVM	0,0103	0,0100	0,0017	0,0000	0,0100	0,0000	0
SVMscaled	0,0100	0,0100	0,0000	0,0000	0,0000	0,0000	0
SVMb	0,0103	0,0100	0,0017	0,0000	0,0100	0,0000	0
SVMbscaled	0,5000	0,5000	0,0000	0,0000	0,0000	0,0000	0
GLM	0,0116	0,0100	0,0051	0,0000	0,0200	0,0000	0
STA13	**0,0168**	**0,0200**	**0,0079**	**0,0100**	**0,0400**	**0,0100**	**0**

Table 4.5.1.1 Electrodes Data.Measures of the performances of the classifiers and of the STA13 scheme. Average figures over 100 iterations

This is a dataset that is particularly suited to classification problems, as can easily be seen from the relatively moderate values for the classifiers' error rate. The best performances are achieved by Quadratic Discriminant Analysis and Linear Discriminant Analysis, a superiority also confirmed by the indicator relative to the best positioning (which indicates how many times the classifier has been the best out of a total of 100 simulations). Given the distribution typology, this result was not unexpected. On the contrary, the error made by the Support Vector Machine in its version "best" with scaling was particularly noticeable and rather curious. The performance of the Stacking, despite the positive component of the parametric classifiers, is quite modest and inclines us not to choose it, as it would be quite onerous with respect to the use of a single classifier with much better performances. In the following box plots relative to the error distribution of the 13 classifiers and of the Stacking in Figure 4.5.1.1, we can appreciate the variability of the error that cannot be captured by the average value, even though the dimensions are much reduced. In this case we can see that there is no variability for several classifiers.

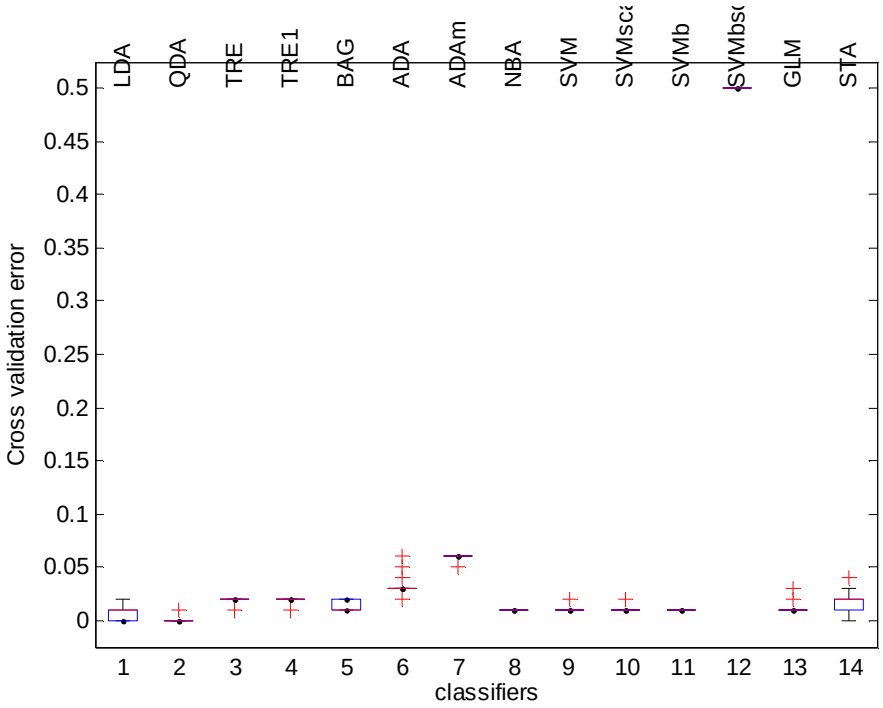

Figure 4.5.1.1. Electrodes Data. Boxplots of distribution error of 13 base classifier and STA 13. Over 100 iterations.

4.5.2 Results for STA6

In this case the model is less complex in terms of the explicative variables (the predictions derived from the base classifiers) that make up the input dataset of the meta-classifier.
Here are the results summarised in the following table 4.5.2.1

	Cross Validation Error	Median Cross Validation Error	Std. Deviation Cross Validation Error	Difference interquartile Cross Validation Error	Range Cross Validation Error	MAD Cross Validation Error	% Best positioning
Classifier							
LDA	0,0047	0,0000	0,0050	0,0100	0,0100	0,0000	62
QDA	0,0020	0,0000	0,0040	0,0000	0,0100	0,0000	38
BAG	0,0148	0,0100	0,0050	0,0100	0,0100	0,0000	0
ADA	0,0293	0,0300	0,0061	0,0000	0,0300	0,0000	0
NBA	0,0100	0,0100	0,0000	0,0000	0,0000	0,0000	0
GLM	0,0127	0,0100	0,0062	0,0000	0,0200	0,0000	0
STA6	**0,0119**	**0,0100**	**0,0061**	**0,0000**	**0,0300**	**0,0000**	0

Table 4.5.2.1. Electrodes Data.Measurement of the performances of the classifiers and of the STA6 scheme. Average values over 100 iterations

Slight differences in the performances of the base classifiers compared to the previous scheme may be due to a different partition in the folds of the cross validation.
Stacking has improved its performance compared to the previous scheme, even though it is not competitive compared to the single classifiers, and it was not even once the best classifier in this case. The superiority of Linear Discriminant Analysis and Quadratic Discriminant Analysis is confirmed, and in any case, parametric base classifiers achieve better results than the others. By means of the following two plots, we can also observe in this case the almost complete absence of variability or oscillations with respect to the repetitions of the error distribution. First of all, the box plot in Figure 4.5.2.1

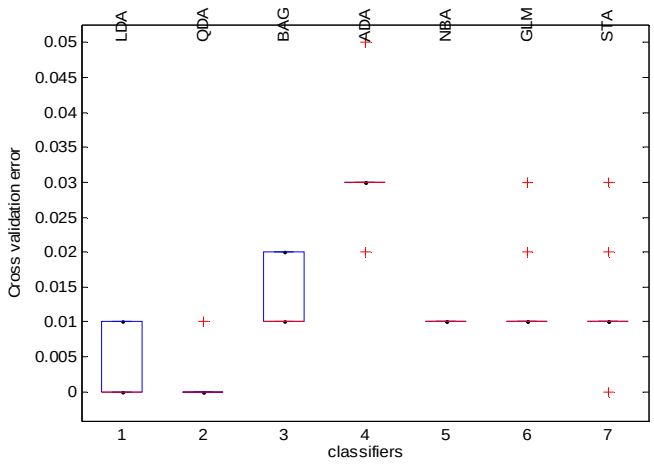

Figure 4.5.2.1. Electrodes Data. Boxplots of distribution error of six base classifier and STA 6. Over 100 iterations.

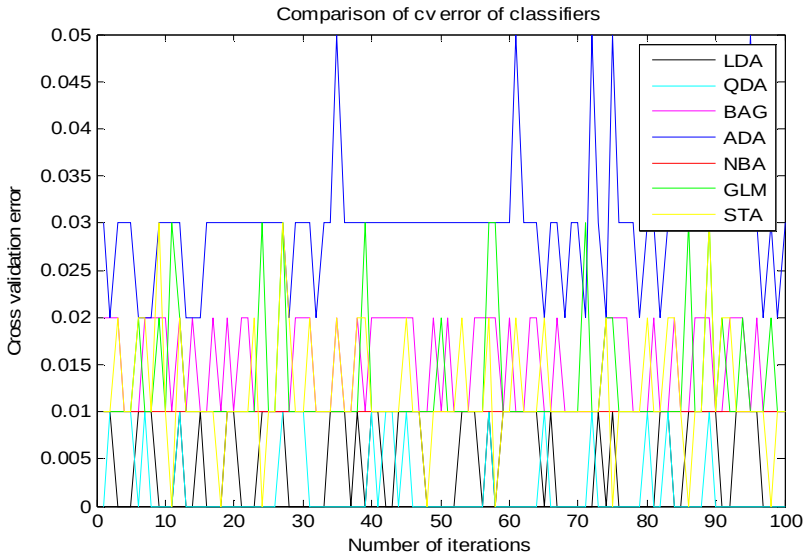

Figure. 4.5.2.2 Electrodes Data. Comparison of the cross validation errors of six base classifiers and STA6 over 100 iterations.

4.5.3. Results for STA3

In this case, as we can see in Table 4.5.3.1, the performance of the Stacking improves further, although it is still never the best in any repetition, but it has come closer to the Support Vector Machine (in the version with the parameter selection), which on average achieves the lowest error and better results that the others which are multiple classifiers.

	Cross Validation Error	Median Cross Validation Error	Std. Deviation Cross Validation Error	Difference interquartile Cross Validation Error	Range Cross Validation Error	MAD Cross Validation Error	% Best positioning
Classifier							
BAG	0,0141	0,0100	0,0049	0,0100	0,0100	0,0000	59
ADA	0,0315	0,0300	0,0081	0,0000	0,0500	0,0000	0
SVMb	0,0100	0,0100	0,0000	0,0000	0,0000	0,0000	41
STA3	**0,0112**	**0,0100**	**0,0033**	**0,0000**	**0,0100**	**0,0000**	**0**

Table 4.5.3.1. Electrodes Data.Measurement of the performances of the classifiers and of the STA3 scheme. Average values over 100 iterations.

It is interesting to note that the Bagged Tree, despite achieving a mediocre performance in terms of average error, is the best classifier for 59 of the 100 repetitions carried out on the procedure.
We can observe that there is not a correspondence for the classifiers between the best positioning achieved and the lowest value of cross validation error achieved.
We omit the box plot which displays, also in this case, an almost complete absence of variability, and we show, for completeness, the graph in Figure 4.5.3.1 that compares the cross validation errors for the 3 classifiers and the Stacking over 100 iterations.

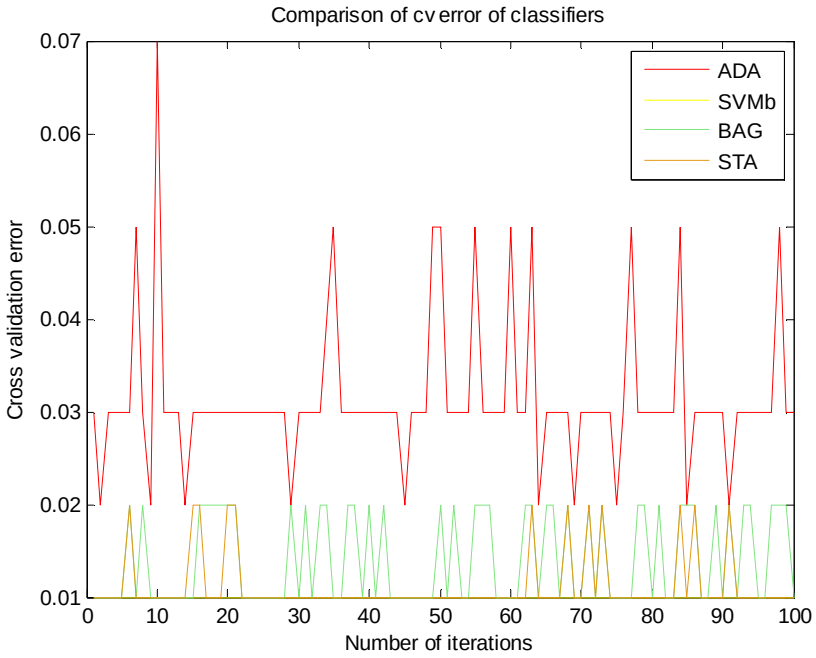

Figure. 4.5.3.1. Electrodes Data. Comparison of the cross validation errors of three base classifiers and STA3 over 100 iterations

4.6 Application to real data: the dataset off-the-book employment

The second application of the Stacking Scheme to real data was carried out on a data sample taken from a dataset containing the results of the inspection surveys carried out by INPS (National Social Security Institute) on Italian companies in order to see if there was any off-the-book employment present. In its original version this dataset is extensive and rather complex, consisting of 14,651 records divided into two non-balanced groups, and 39 variables that are both continuous and qualitative.
The data used in the analysis are taken from 230 records extracted from the original dataset, respecting, where possible, the proportions of the data originally present in each group, and we have selected 4 variables.
Furthermore, because of the marked level of overlap between the two groups and the presence of collinearity between variables, these variables have been transformed by taking logs.

n_1 =126

n_2 =104

Identification of the variables used:

Org4 = No. paid days of unskilled workers per unit of total paid day
Dim8 = Total paid days
Pers13 = employee expenses per employee (Asia)
Pers16 = Productivity per employee
Label class = Absence/Presence of the off-the-book workforce during the last INPS inspection

Three different Stacking schemes were used for this case too:

Scheme 1) 3 base classifiers with 4 explanatory variables + LR/RR as mutually exclusive meta-classifiers
Scheme 2) 6 base classifiers with 4 explanatory variables + LR/RR as mutually exclusive meta-classifiers
Scheme 3) 3 base classifiers with 4 explanatory variables + LR/RR as mutually exclusive meta-classifiers

Since the dataset does not come from our simulation plan but refers instead to real data, and since there has also been a transformation of the original variables, it is interesting to observe Figure 4.6.1 which shows the scatter plot matrix.

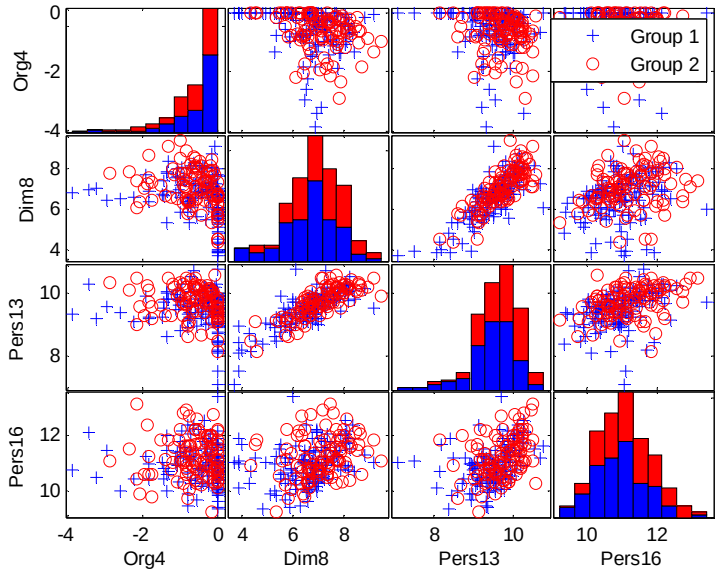

Figure 4.6.1. Off-the-book employment Data. Scatterplot matrix with bivariate scatters of the three variables and histograms on the main diagonal. Units in Group 1 are represented by blue crosses

Both the conditions of extensive collinearity between the variables and of the degree of overlap between the groups seem to have improved compared to the original situation, although the overlap remains. There do not seem to be any variables that make a perfect separation possible between the two groups.
We now move on to the analysis of the main results for each proposed Stacking Scheme:

1) Scheme STA13

In the scheme made up of 13 base classifiers, as shown in Table 4.6.1 below, the average cross validation error stays quite high for all the classifiers, as was predictable given the level of overlap between the groups.

	Cross Validation Error	Median Cross Validation Error	Std. Deviation Cross Validation Error	Interquartile Difference Cross Validation Error	Range Cross Validation Error	MAD Cross Validation Error	% Best positioning
Classifier							
LDA	0,4347	0,4348	0,0086	0,0087	0,0348	0,0043	2
QDA	0,4653	0,4652	0,0125	0,0174	0,0652	0,0087	0
TRE	0,4701	0,4696	0,0288	0,0413	0,1609	0,0217	0
TRE1	0,4642	0,4696	0,0307	0,0457	0,1435	0,0261	5
BAG	0,4772	0,4783	0,0224	0,0261	0,1087	0,0130	0
ADA	0,4360	0,4348	0,0192	0,0261	0,0913	0,0130	14
ADAm	0,4395	0,4348	0,0220	0,0348	0,1087	0,0152	16
NBA	0,4351	0,4348	0,0093	0,0087	0,0522	0,0043	2
SVM	0,4240	0,4217	0,0110	0,0130	0,0609	0,0087	17
SVMscaled	0,4236	0,4217	0,0129	0,0130	0,0565	0,0087	21
SVMb	0,4240	0,4217	0,0110	0,0130	0,0609	0,0087	0
SVMbscaled	0,4299	0,4304	0,0132	0,0217	0,0609	0,0087	9
GLM	0,4318	0,4304	0,0081	0,0130	0,0391	0,0043	0
STA	**0,4417**	**0,4435**	**0,0286**	**0,0391**	**0,1391**	**0,0217**	**14**

Table 4.6.1 Off-the-book employment Data . Measurements of the performances of the base classifiers and the Stacking scheme STA13 calculated with reference to the respective distribution of the cross validation errors. Average values for 100 iterations.

Having chosen to carry out the evaluation and comparison of the classifiers by estimating cross validation (medium in this case), but also with the aim of broadening the representation of the error distribution, we included the calculation of other indicators which could improve our knowledge and the plots shown above.
When comparing the performance achieved by Stacking with those of some base classifiers (Linear Discriminant Analysis, Adaboost, Support Vector Machine, and Logistic Regression) it does not seem to be competitive compared to the use of a single classifier. However, it should be noted that, although Stacking has quite a modest result in terms of accuracy, it manages to achieve quite an interesting result in terms of best positioning (it was the best in 14% of the iterations) which is higher than classifiers characterised by lower average values for cross validation error. With regard to this, Figure 4.6.2 and Figure 4.6.3 show the presence of a certain variability in the error distribution for some classifiers on the total of the iterations.

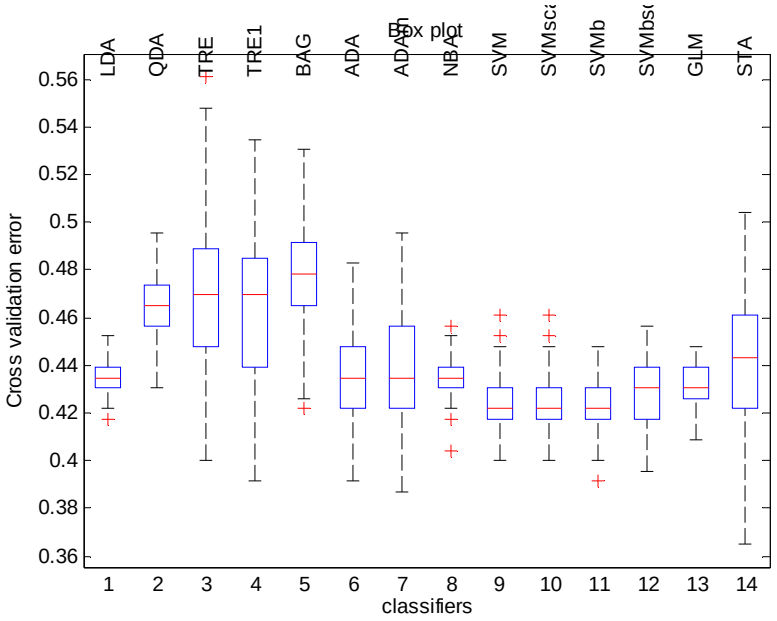

Figure 4.6.2. Off-the-book employment Data. Boxplots of error distribution of 13 base classifiers and Stacking scheme. Over 100 iterations.

In this sense, examples are represented by Stacking and by the Classifcation Tree, which achieve respectively the minimum and maximum values (compared to the other classifiers) of cross validation error in some iterations.

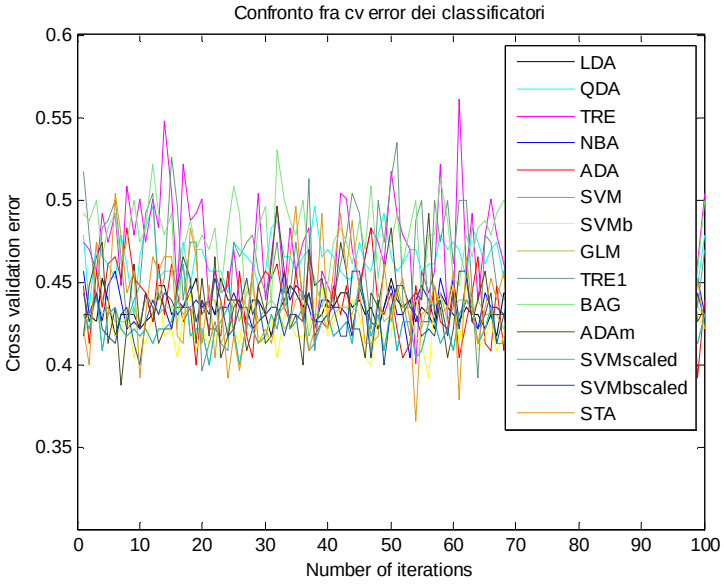

Figure 4.6.3. Off-the-book employment Data. Comparison of cross-validation error of 13 base classifiers and Stacking scheme. Over 100 iterations.

2) Scheme STA6

In the scheme with 6 base classifiers, the average cross validation error for Stacking is lower than any other single classifier and its performance is the best in terms of the number of times when it was the best classifier for the total of iterations carried out, as we can see in Table 4.6.2 below.

Classifier	Cross Validation Error	Median Cross Validation Error	Std. Deviation Cross Validation Error	Interquartile Difference Cross Validation Error	Range Cross Validation Error	MAD Cross Validation Error	% Best positioning
LDA	0,4342	0,4348	0,0091	0,0130	0,0565	0,0043	16
QDA	0,4609	0,4609	0,0153	0,0174	0,1043	0,0087	2
BAG	0,4704	0,4696	0,0221	0,0261	0,1435	0,0130	2
ADA	0,4361	0,4348	0,0184	0,0217	0,1217	0,0130	22
NBA	0,4336	0,4348	0,0094	0,0130	0,0696	0,0043	15
GLM	0,4331	0,4348	0,0095	0,0130	0,0565	0,0087	8
STA	**0,4309**	**0,4304**	**0,0243**	**0,0304**	**0,1478**	**0,0174**	35

Table 4.6.2. Off-the-book employment Data. Measurements of the performances of six base classifiers and the Stacking scheme STA6, calculated with reference to the respective distribution of the cross validation errors. Average values for 100 iterations.

It would seem, therefore, that this combination of classifiers is the one that best expresses the predictive capacities of Stacking in terms of accuracy. In this case too, as shown in Figure 4.6.4, Stacking is characterised by a certain variability compared to the others, even though the average and median values for cross validation error are the lowest.

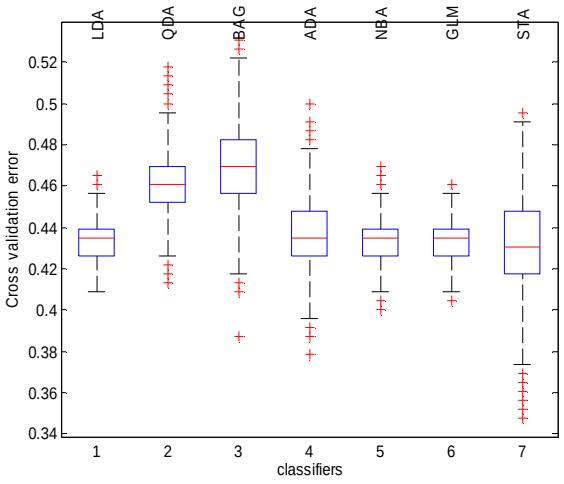

Figure 4.6.4. Off-the-book employment Data. Boxplots of error distribution of six base classifiers and STA6 scheme. Over 100 iterations.

The fluctuations in cross validation error compared to some classifiers for the total number of iterations, especially for Stacking and Bagging, are shown in Figure 4.6.5. This variability is, of

course, characterised by the fact of achieving minimum values for the former and maximum values for the latter, compared with the other classifiers which, on the contrary, appear to have quite moderate fluctuations.

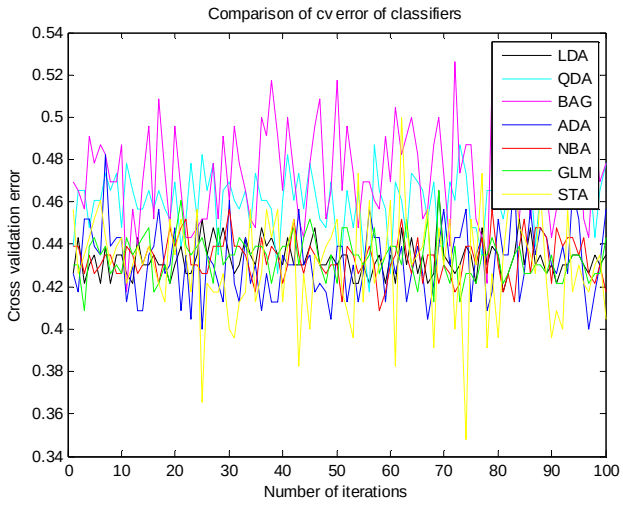

Figure 4.6.5. Off-the-book employment Data. Comparison of cross-validation error of six base classifiers and STA6 scheme. Over 100 iterations.

3) STA3

In the scheme with three base classifiers, Stacking seems competitive when compared to the use of ensemble methods (Bagging and Adaboost), but not preferable to the use of the Support Vector Machine which has a lower error rate in the classification.

	Cross Validation Error	Median Cross Validation Error	Std. Deviation Cross Validation Error	Interquartile Difference Cross Validation Error	Range Cross Validation Error	MAD Cross Validation Error	% Best positioning
Classifier							
BAG	0,4667	0,4652	0,0252	0,0304	0,1304	0,0174	1
ADA	0,4352	0,4348	0,0207	0,0261	0,0957	0,0130	28
SVMb	0,4243	0,4261	0,0138	0,0174	0,0783	0,0087	51
STA	**0,4311**	**0,4304**	**0,0178**	**0,0174**	**0,1000**	**0,0087**	**20**

Table 4.6.3. Off-the-book employment Data. Measurements of the performances of three base classifiers and the Stacking scheme STA3, calculated with reference to the respective distribution of the cross validation errors. Average values for 100 iterations.

The variability of Stacking is more moderate compared to what we have seen in previous schemes and compared to ensemble methods, as illustrated in Figure 4.6.6. and Figure 4.6.7.

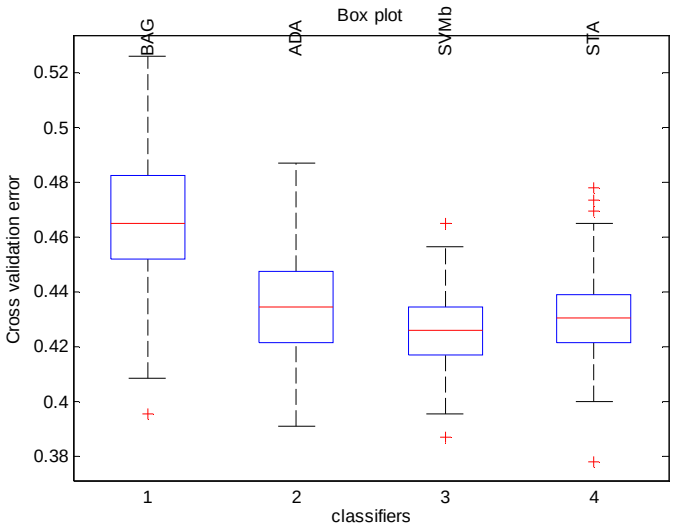

Figure 4.6.6. Off-the-book employment Data. Boxplots of error distribution of three base classifiers and STA3 scheme. Over 100 iterations

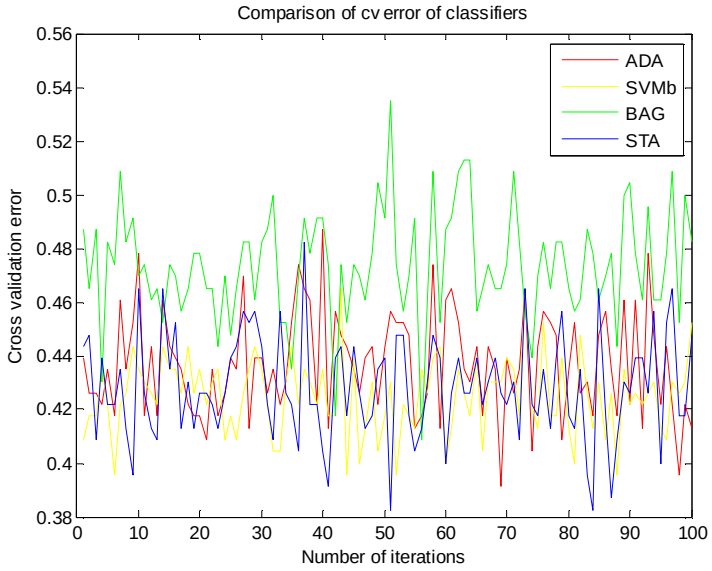

Figure 4.6.7. Off-the-book employment Data. Comparison of cross-validation error of three base classifiers and STA3 scheme. Over 100 iterations.

If we summarise some of the main results obtained in terms of the Stacking scheme's performance, using subsets of base classifiers of different sizes and typologies, for real datasets and simulated data, we can then analyse the following Table 4.6.4.

	STA3	STA6	STA13
Input Data 0-level			
120_3_05	0,3547	**0,3472**	0,3545
120_10_05	**0,2323**	0,2399	0,2467
120_3_3	**0,0048**	0,0066	0,0073
120_10_3	**0**	**0**	0,0001
200_3_05	0,3498	**0,3450**	0,3475
200_10_05	**0,2268**	0,2324	0,2314
200_3_2	**0,0434**	0,0462	0,0465
200_10_2	**0,0004**	0,0016	0,0016
200_3_3	0,0058	**0,0016**	0,0070
200_10_3	**0**	**0**	**0**
Electrodes	**0,0112**	0,0119	0,0168
Off-the-book	0,4311	**0,4309**	0,4417

Table 4.6.4. Off-the-book employment Data. Cross validation average error rate for different Stacking schemes and different base-level datasets.

There would seem to be a confirmation of the circumstance that a greater complexity of the meta model does not improve results: in fact the best Stacking performances were achieved (at least in the examples analysed) with schemes with a lower number of base classifiers STA3 and STA6. We did not take the STA4 scheme into consideration, because with respect to a low second level complexity, it recorded very bad performances due to the weight of the "weak" component among its classifiers.
STA13 achieves the same level of performances as the other two schemes only with the hypothesis that both the degree of complexity and the degree of separation between the groups are at a maximum in the base-level input datasets.

4.7 Discussion

In this chapter we have summarized the main results obtained from the application of the proposed Stacking scheme to datasets generated by means of the experimental design and also real data. We were interested in investigating, empirically, the performances of the single classifiers and of Stacking for datasets with different characteristics, and above all if it is convenient to use Stacking in terms of improving performances instead of a single classifier, bearing in mind the necessary increase in computation. An analysis of the results obtained by applying the Stacking scheme to the set of the datasets generated by means of the simulation study, does not lead us to believe that the Stacking scheme is preferable in terms of performances to the use of the best single classifier. It always achieves good performances and is to be considered among the best, but it does not seem to be preferable for this type of application. In the case of contaminated data, Stacking improves its performances noticeably compared to what we have observed for non-contaminated data, in some cases also in comparison with scaled SVM, and generally appears to be very competitive, above all when the contaminations are more substantial. In a set of classifiers in which there was no scaledSVM it

would be the best for each of the analysed datasets. The results obtained using three different Stacking variants with different base-level classifier subsets built on different datasets show that the choice from the start of high performance base-level classifiers could increase the predictive capacities of Stacking, and however, a comparison shows that the Stacking scheme with three base classifiers, is more successful (although the differences are slight). There would seem to be a confirmation of the circumstance that a greater complexity of the meta model does not improve results also for the applications to real data : in fact the best Stacking performances were achieved (at least in the examples analysed) with schemes with a lower number of base classifiers STA3 and STA6.
We can observe, moreover, that there is not always a correspondence for the classifiers between the best positioning achieved and the lowest value of cross validation error achieved and also because, since we are dealing with an average of 100 iterations, the variability of a classifier is significant in terms of the error returned which is often very high. It may therefore happen that classifiers with a higher variability can achieve better positioning. Because of the variability and fluctuation of the cross validation error, the average does not seem to be sufficient as a measurement and this circumstance should lead to us to continue looking for more adequate measurements (at some point combining the use of more than one index) which are able to capture accuracy in the best possible way in terms of estimating the prediction error returned by the single classifiers also in order to improve the comparison with the use of a more complex scheme like Stacking.

CHAPTER 5

Forward Search in multivariate analysis in the context of supervised classification methods

5.1 Introduction

Forward Search (Atkinson, Riani and Cerioli (2004)), is a methodological proposal which, apart from allowing anomalous values to be identified, also makes it possible to monitor in an iterative way the effect exerted by each unit on the model and on the quantities of interest in each step of the algorithm.

Starting from the construction, using robust methods, of a subset $S(m)$ free from anomalous values and which represents the heart of the distribution, a dynamic implementation will be achieved, increasing the size of the robust sample selected with the introduction of one observation at a time. The choice of the new cardinality subset *m+1* is found using the Mahalanobis distances calculated to the step *m.* More precisely, the observations are chosen with the *m+1* with the smallest distances to form the new subset *S(m+1).* The process is repeated at every step of the search and continues until $m = n$.

The "philosophy" at the heart of the Forward Search approach is the creation of a dynamic data analysis process.

The innovative nature of the proposal for this part of the thesis is to be found chiefly in the extention of the Forward Search approach to the Stacking scheme in order to build the procedure in a robust way and to monitor the effects that each observation, outlier or not, can exert on the model and on the quantities of interest.

A brief and general illustration of the main phases of the Forward Search process in a multivariate context will be given at the beginning of this chapter, followed by a description of the procedure for constructing the Stacking scheme which is then inserted into a Forward Search framework.

5.2 The Forward Search

5.2.1 Starting the Search

We find an initial subset of moderate size by robust analysis of the matrix of bivariate scatterplots. The initial subset of r observations, which we denote with Sr^*, consists of those observations which are not outlying on any scatterplot, found as the intersection of all points lying within a robust contour containing a specified portion of the data (Riani and Zani 1997) and inside the univariate boxplot. There are two versions of the robust bivariate contour. The first uses convex hull peeling and *B-spline* smoothing (Zani, Riani and Corbellini 1998). The

second, computationally more simple, in which ellipses with a robust centroid are fitted to the data (Riani and Zani 1997).

5.2.2 Progressing in the Search.

At each step of the forward search, given a subset $S_*^{(m)}$, we estimate the parameters μ and Σ, we compute and order all squares Mahalanobis distances d_{im}^{*2}, with $i = 1,...,n$ for $m_0 \leq m \leq n$ and we select the (*m+1*) observations with the (*m+1)* smallest distances to form the new subset $S_*^{(m+1)}$. Here *m* runs from m_0 to the fit to all observations when $m = n$. Usually one observation is added at a time, but the inclusion of an outlier can cause the ordering of the observations to change. An interchange occurs and two or more new units enter the subset in a particular step. Of course, at least one unit then has to leave the subset in order for the size to increase by one unit.

5.2.3 Monitoring the Search.

Outliers and influential observations can be detected graphically monitoring particular Mahalanobis distances:

- **Minimum Mahalanobis distance among the units which do not belong to the subset**
  $d_{[m+1]}^2 \quad m = m_0,...,n-1$
- **Maximum Mahalanobis distance among the units inside the subset**
  $d_{[m]}^2 \quad m = m_0,...,n$

If there are outliers they will have large distances during the early part of the search that decrease dramatically at the end as the outlying observations are included in the subset of observations used for parameter estimation. If our interest is in outlier detection we can also monitor, for example, the minimum Mahalanobis distance among units not in the subset. If an outlier is about to enter, this distance will be large, and the graph will show a peak in the relative curve in the previous step to that of the inclusion of the outlier, although it will decrease again as the search progresses if a cluster of outliers join.
On the contrary, the plot monitoring the maximum Mahalanobis distances for units inside will show a peak when the first outlier enter the subset.

5.3 Extending Forward Search to Stacking schemes

As we have already seen in the part of the thesis dedicated to the traditional approach, the entire process of the proposed Stacking scheme was implemented in a Matlab environment in this case too. We created a specific and suitable routine which performs Forward Search in multivariate analysis in the context of the combination of supervised classification methods, which is then inserted into the field of the FSDA toolbox, created for multivariate data analysis.
This procedure is, in any case, more complex than the other one, because as well as the construction of the Stacking scheme, there is also its inclusion into a Forward Search context to be considered, which in turn has been extended to the field of supervised classification. Thus

the typical Forward Search procedures are extended to Stacking, procedures which refer to the choice of the best robust subset, the criteria for search progress and the creation of specific plots that support the monitoring of the quantity of interest graphically.
The construction of single classification models and of the final classifier will be partly different compared to what we have seen in the traditional approach, since the particular dynamic characteristics linked to Forward Search must be taken into consideration. As far as the general description of the base classifier characteristics used is concerned, this, together with Stacking, the construction of homogenisation procedures and the generation of outputs are referred to in detail in Chapter 3.
Initially a subset will be selected with robust methods - in which ellipses with a robust centroid are fitted to the data - for each group separately, on which the fit of the single algorithms will be carried out for the construction of the base classifiers, with the certainty that no outliers are present in the subset.
Only "good" units for each population are chosen for the fit of the model, which will therefore be constructed in a "robust" way. Unlike the traditional approach, here we start from the robust subset, which is the heart of the distribution, which will be implemented iteratively during the search and on which, at every step, the entire Stacking scheme process will be repeated until $m = n$. This will allow us to see a film of the data (Atkinson, Riani and Cerioli (2004)), in contrast with the static vision of the traditional approach, so as to be able to appreciate the effect that each unit has on the "fitted" model and on the quantities that interest us. The strength of the search in this case lies in the graphic representation and the plots generated by the procedure which will allow us to monitor several interesting aspects, especially in the presence of outliers. In particular we can observe the effects on the performances of the classifiers in terms of prediction error and on the rule of decision.
Since we are dealing with classification problems, within the subset the units are subdivided into groups, two in our case, which will have to be big enough for the estimation of the model and its parameters for the computation of the Mahalanobis distances. Generally speaking we try to respect the proportion of the units present in the groups of the original dataset.

With regard to this, before describing the procedure, let us introduce two fundamental aspects that concern:

- The inclusion criteria of the units in the subset
- The parameters of the model for calculating the Mahalanobis distances

As far as the inclusion of the units in the subset is concerned, it can happen in two ways:

- **Balanced**. The units will be inserted into the subset, taking into account the ordering of the Mahalanobis distances carried out separately for each group, respecting where possible the existing proportion between the groups in the original data.
- **Not balanced**. The ordering of the units will be carried out without considering what belongs to which group, and the inclusion will take place following the general criterion of the *(m+1)* units with the smallest Mahalanobis distances. In this hypothesis, if outliers should be present only in one group, they will make their entrance in the final part of the search.

As far as the aspects regarding the parameters for calculating the Mahalanobis distances are concerned, it will be possible to:

- Estimate the centroid μ_ℓ from the subset and the matrix of variance and covariance for each group Σ_ℓ and, in such a case the distances will be:

$$d^2_{i\,(m\ell)} = (y_i - \hat{\mu}_{(m\ell)})^T \hat{\Sigma}^{-1}_{(m\ell)} (y_i - \hat{\mu}_{(m\ell)})$$

- Estimate the centroid μ_ℓ from the subset for each group and the matrix of variance and pooled covariance Σ_W in the hypothesis in which the matrices are the same for both groups. The distances will be: $d^2_{i\,(m\ell)} = (y_i - \hat{\mu}_{(m\ell)})^T \hat{\Sigma}^{-1}_W (y_i - \hat{\mu}_{(m\ell)})$

In the part dedicated to the applications to data we will see that these situations will provide different and very interesting results.
Generally speaking, it does not seem advisable to indicate one setting rather than another a priori.
In our applications, a balanced search will be used first and the calculation of the distances will be made based on the estimation of μ_ℓ and Σ_W (centroid of each group and matrix of variance and pooled covariance, common for both groups). However, since the Stacking scheme includes the use of different classification methods, besides Discriminant Linear Analysis, so the hypothesis of the existence of a pooled covariance matrix is not required. Instead the use of a specific covariance matrix for each group will make it very interesting to explore various possibilities, as we will also in the hypothesis of a non-balanced search.
We will see in the procedure description that some choices have been made for Stacking that are partially different from those made for the traditional part of the analysis, in order to facilitate its insertion into the Forward Search context.

At each step of the search:

- The fit will be carried out on the model for the units of the robustly chosen subset, together with the estimation of the parameters for calculating the Mahalanobis distances, whose ordering will be carried out separately for each group in the hypothesis of the choice of a balanced search, or, indiscriminately, in the hypothesis of a non-balanced search. Each classification algorithm will then be trained on the subset and, iteratively, with the insertion at each step of a new unit (chosen with the criteria mentioned above), the different base classifiers will be fitted, then, on the training data, until the dimension of the subset is equal to n.
  The algorithms in this phase will also be used in a particular predictive function on the units of the entire dataset *n*, and will give back at every step the class predicted for each unit of *n*, and the a posteriori probabilities relative to each unit of *n* which, however, will have to be considered as *resubstitution* posterior probabilities (or "in sample") for that part of the unit *m* used for the fit of the model and will belong to the subset, and as *cross-validation* posterior probabilities, which are extensively intended, for that part of unit (*n-m)*, which, not being part of the subset, haven not been part of the fit of the model. The predictions relating to the class of belonging for the (*n-m*) units that do not belong to the subset will make an estimation of the cross validation error possible for the units outside of the subset.
- A ten-fold stratified cross-validation procedure will also be applied to the subset data, a procedure for all the classification models generated robustly, in which the same data

partition is used for all the methods, in order to make outputs homogenous and perfectly comparable. At each step, by means of the procedure, the predictions will be given back that will allow us to estimate the cross validation ("in sample") error relative to each classifier for the units inside the subset as well as the probabilities of belonging to the class which, in this case, represent probabilities of cross validation ("in sample") error, always for the units that belong to the subset. These units will represent the input dataset for the meta-classifier, or rather, in this case too, will represent the explicative variables for linear regression and the ridge regressions that will represent the combiner functions in a mutually exclusive way as described in the traditional part (where hypotheses of multicollinearity recur).

- The fit of the meta-classifier will be carried out on the a posteriori probabilities relative to each base classifier generated by the cross-validation procedure on the subset data. Furthermore the meta-level classifier, used as a predictive function, will give back the predictions relative to the Stacking scheme which will make possible an estimation of the cross-validation error and the a posteriori probabilities of cross-validation for Stacking for the units inside the subset.
  The cross-validation error is a reliable estimate of prediction error, both of base classifiers and of the Stacking scheme, since the data for prediction are not the same as those used for the fit of the algorithms and therefore for constructing the decision rule. It will be used in this case too, as in the traditional approach, for the evaluation of base classifiers and of Stacking, but its graph will be able to supply us with indications about the behaviour of the decision rule, by monitoring the effect caused by the entrance of possible outliers to the cross- validation error.
- The fit of the meta-classifier will also be carried out by using, instead of the probability of cross validation ("in sample"), the posteriori probabilities of resubstitution ("in sample") given by the base classifiers in the predictive phase relative to that part of the unit *m* that took part in the fit of the model. Similarly to what was carried out for base classifiers, the meta-classifier, used as a predictive function on the units of the entire dataset *n*, will give back at every step the class predicted for each unit of *n*, and the a posteriori probabilities relative to each unit of *n*, which must be considered as resubstitution posterior probabilities (or "in sample") for that part of unit *m* used for the fit of the model and which therefore belong to the subset, and as cross-validation posterior probabilities, extensively intended, for that part of the unit (*n-m)*, which, not being part of the subset, did not take part in the fit of the model.
  In this case, the predictions relative to the class of belonging for the (*n-m*) units that do not belong to the subset will make possible the estimation of the cross validation error for the units outside of the subset and the posteriori probabilities of cross validation given back by this Stacking scheme must be seen as being relative to the units that are outside the subset.

At the end of the Stacking process implemented in a Forward Search context, when all the units have entered the subset, the procedure returns numerous plots that are a fundamental tool for identifying any anomalous values, for exploring data and for monitoring the performances of the various base classifiers and the final classifier, as well as the behaviour of the decision rule that presides over the functioning of each algorithm and of Stacking.

The graphs produced by the procedure enable us to monitor at every step of the search the following:

- **Minimum Mahalanobis distance among the units which do not belong to the subset and (m+1)ordered MD.** This is the most important graphic tool, together with the following one which deals with the maximum, for the exploration of the data and, especially for the detection of anomalous values present in one or both groups. If one (or more) anomalous observation enters a group, the graph will show a peak in the relative curve in the previous step to that of the inclusion of the outlier.
- **Maximum Mahalanobis distance among the units inside the subset and m-th ordered MD.** On the contrary, in this case, if one (or more) anomalous observation enters a group the graph will show a peak in the corresponding curve one step later of the inclusion of the first outlier. Generally speaking, when an outlier enters the subset, the size of the peak in the plot of maximum distances is smaller than that in the plot of minimum. The characteristics of the minimum and maximum curves are closely linked to the choice between a balanced or non-balanced search. In the case of a non-balanced search, if there are outliers in a group, they will enter the last steps of the search. In the case of a balanced search, the entrance might occur in the central part of the search.
- **Posterior probabilities.** This is the graph for the a posteriori probabilities of the *n* units, which, as indicated, will be of cross validation for the units that did not take part in the fit of the model, in the estimation of the μ_ℓ and Σ_ℓ parameters, and of resubstitution for the units that did take part in the fit. This is a useful tool for data exploration.
- **Cross-validation posterior probabilities.** This graph shows the posterior probabilities deriving from the cross validation procedure for the subset units, and those of cross validation extensively intended for that part of the unit *(n-m)* that is not part of the subset and thus did not take part in the fit of the model. It is a useful tool for data exploration.
- **Cross-validation error estimate for the units inside the subset.** In combination with the graph of the minimum Mahalanobis distance, it can be a valid tool for analysing the effect of the entrance of anomalous values on the decision rule, by monitoring how the error is modified after their entrance. The objective is in fact to analyse the effects that may be caused by modifications to the input dataset on the performances of the single classifiers and of Stacking, and, based on these performances, to evaluate whether it is worthwhile using the Stacking scheme or, more sparingly, a single classifier.
- **Cross-validation error estimate for the units outside the subset.** In combination with the plot of the cross-validation error for the units inside, it can be a valid tool for analysing the effect of the entrance of anomalous values, since if the internal and the external subset are rather homogeneous we expect that the two errors (in or out the subset) are the same. On the contrary, larger errors for units not belonging to the subset indicate that, in the next steps, units not in accordance with to those are the subset will be included.

In the traditional approach, for the hypothesis in which there are outliers in the input dataset for the base classifiers, the model is also fitted to the outliers. In this case the predictions derived, the a posteriori probabilities and, obviously, the error estimations could be incorrect. In the particular case of Stacking, which uses a posteriori probabilities deriving from the base classifiers as an input dataset, the non-correctness of the estimations could also be reflected in the output of the final classifier. Thus the performance of the Stacking could also be degraded.

The presence of outliers in the input dataset may therefore be reflected in the predictions of the final classifier.

5.4 Discussion

In the passage from the traditional approach to Forward Search, we acquire a much wider knowledge of the structure of the data, the trajectories of the observations, the importance exerted by each of these factors, whether outlier or not, on the model or on the aspects that are of interest for evaluating the performances of the Stacking scheme. With the traditional modality we could only see a static result, which is equivalent to the last step of the search but no information about the effect exerted by the single units, since they were treated simultaneously in the analysis, even if there was a logic of sample rotation such as cross-validation. In other words, it was not possible to appreciate the entrance of anomalous values or to see at which point the classification rule may become unstable. On the other hand, being able to identify the entrance of anomalous values from the graph of the minimum Mahalanobis distance allows us to make decisions and in any case to evaluate classification error and therefore the stability of the decision rule up until the moment when the outliers enter, or rather until the data that take part in the fit of the model are altered, with the consequence that the predictions generated by the model are also altered. Thus the user, by identifying the entrance of anomalous values, will be able to decide whether or not to eliminate from the dataset those units that are not in accordance with the model with the aim of improving its performances.

CHAPTER 6

Experimental results

6.1 Introduction

This chapter illustrates the results of some applications of the Stacking scheme proposed in a Forward Search context. For this analysis we used both datasets generated by the experimental plan as seen in Chapter 4 and some real datasets. The same Stacking scheme with 7 base classifiers was used in each application, with linear regression and ridge regression as meta-classifiers. The general formulations of four proposed algorithms are summarised in Appendix to Chapter 3, while in section 3.3 our specific implementations carried out in a Matlab environment and relative to each algorithm are illustrated. First the results of the two experiments on simulated data will be displayed and then in the following sections space will be given over to the application of the procedure to real data.

6.2 Simulated data

The first dataset used comes from the experimental plan. We have two groups of 60 observations generated from a standard multivariate normal population.

n_1 = 60

n_2 = 60

v = 3

δ = 1.5

The First Group consists of a 60x3 matrix generated from a standard multivariate normal population with a mean equal to 0 for all variables and a covariance matrix Σ.
The Second Group was also generated from a multivariate normal population, but with a mean equal to 1.5.
We therefore imposed a slight degree of separation between the groups.
The analysis was carried out for all the datasets (simulated and real) with a Stacking scheme that includes the use of 7 base classifiers:

- Linear Discriminant Analysis (LDA)
- Quadratic Discriminant analysis (QDA)

- Classifcation Trees (TRE)
- Bagged Classification Trees (BAG)
- AdaBoostM1 (ADA)
- Naive Bayes (NBA)
- Logistic Regression (GLM)

As well as the use of Linear Regression and Ridge Regression with mutual exclusion as the meta-level combiner, or rather with the intervention of Ridge hypotheses of multicollinearity occur.
First, we analyse the scatterplot matrix in Figure 6.2.1 of the three variables plotted with different symbols for each group. Although we have simulated the data, so we know them quite well, the scatterplot matrix still represents an interesting starting point for analysis because it shows much of the structure of the data.

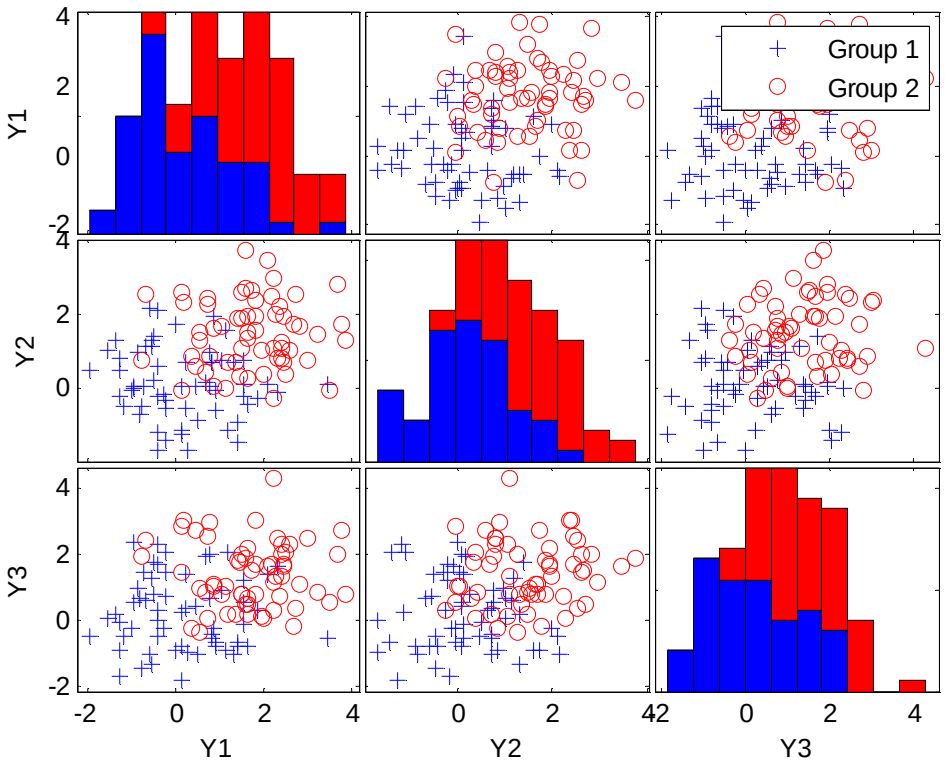

Figure 6.2.1. Simulated data. Scatterplot matrix with bivariate scatters of the three variables and histograms on the main diagonal. Units in Group 1 are represented by blue crosses.

The scatterplot matrix shows a slight degree of overlapping among the two groups for all pairwise scatter plots of the variables. None of the two groups is perfectly separated in any of the plots and the form of the bivariate distribution seems quite similar in all the groups. We run the Search by performing the whole Stacking scheme on the units included in the robust subset at each step. We start with a subset size of 48 out of the total 120 units, and a balanced (constrained) search and a common covariance matrix for the two groups. The subset will contain an equal number of observations from each group and the search will usually proceed with the insertion of one unit for each group.
The search can be used to explore the structure of the data since it progressively includes units with small Mahalanobis distances. Thus, at the beginning of the search there will be the units that are most in agreement with the model to be fitted. The purpose of the search may therefore be not only to identify atypical observations, but also to determine which effect they have on the model to be fitted and to monitor the quantities of interest.
The plots of the minimum Mahalanobis distance of the units outside the subset in Figure 6.2.2 and the maximum Mahalanobis distance of units belonging to the subset are useful tools for detecting the presence of outliers in a group, and also for exploring the data. The search progressively includes units with small Mahalanobis distances: if the distance is larger the outlier is introduced into the subset. From our observation of the plots, it seems no outlier is present.

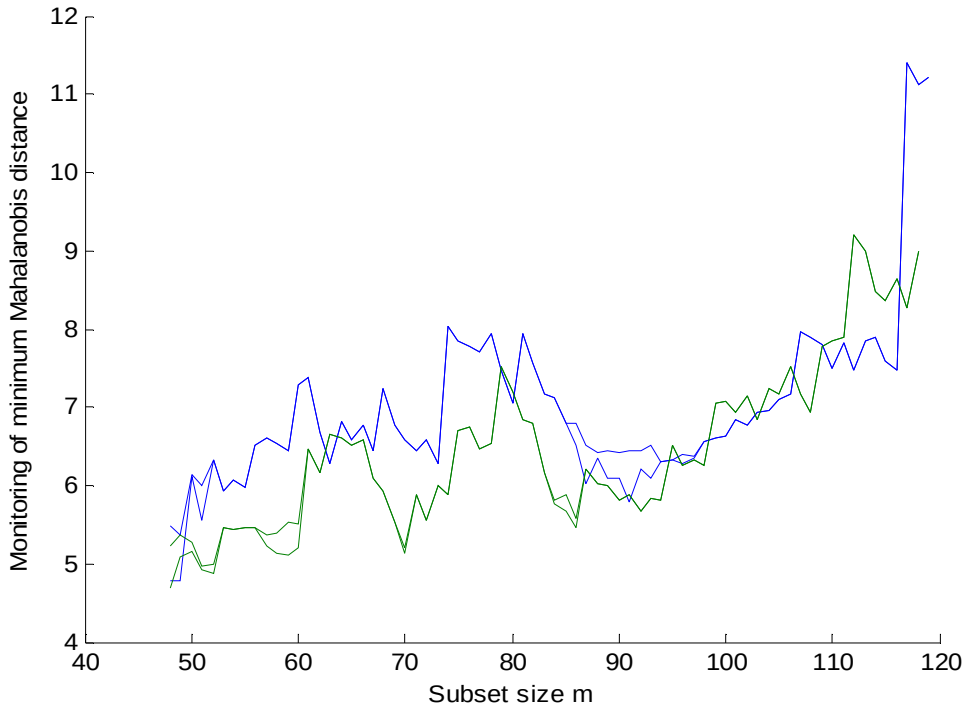

Figure 6.2.2. Simulated data. Plot of minimum Mahalanobis distances for units not belonging to the subset. Blu line, Group 1. Balanced search.

In the initial phase of the search it seems that there is a larger distance for Group 1. The appearance of the curves of the Mahalanobis distances may depend on the choice made between a balanced and an unbalanced search.
We are interested in the evaluation of the classifiers and of Stacking at each step of the search and in monitoring the stability of the decision rule and, therefore, the effect that each unit has on the probabilities for monitoring the decision rule. It seems quite clear from the plots of the posteriors and of the cross validation posterior probabilities that the classifiers display different behaviour also in this case where there are no outliers. We are interested in the performance of the classifiers and therefore in the stability of the decision rule and of the changes that may intervene on it at the entrance of some units that are not necessarily anomalous values, as it seems in this example. It should be remembered that in the posterior probabilities there are also those relating to the units that took part in the fit of the model and also in the estimation of the parameters for the computation of Mahalanobis distances, in this case with a pooled covariance matrix, while the cross validation posteriors never refer to units that have taken part in the fit of the model. For monitoring the decision rule we monitor the cross validation errors for the units inside in Figure 6.2.3 and outside the subset at each step of the search. There is quite a moderate error for many classifiers, and at the beginning of the search (for the units that are most in agreement with the model) it is close to 0 for some classifiers, while for the Classification Tree we notice a slightly higher error level in the central part of the search which is moderated in the last steps.

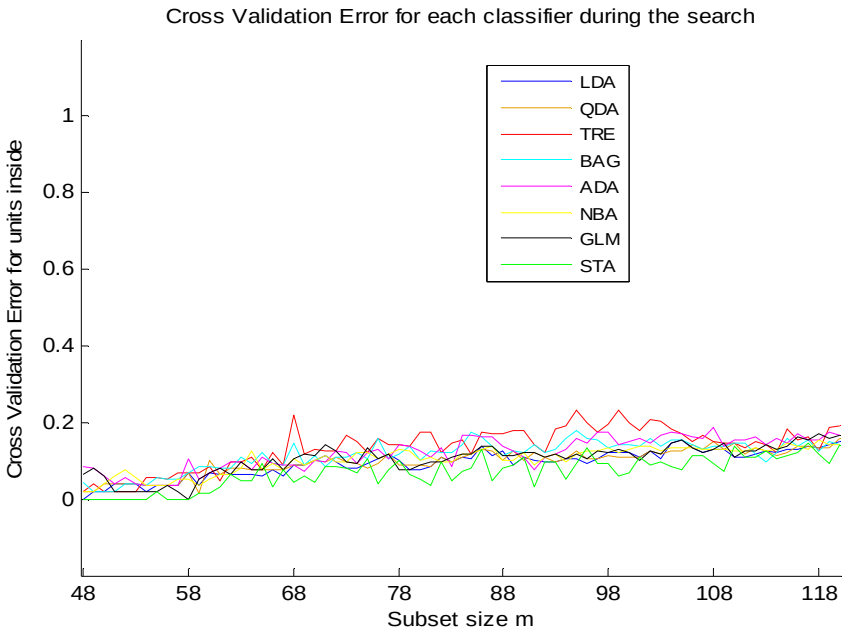

Figure 6.2.3. Simulated data. Cross-validation error for units belonging to the subset of the seven base classifiers and Stacking. Balanced search.

Stacking seems to be preferable in terms of cross validation error and is the best or among the best at every step of the search. This seems to continue to be the case even if we change the choice of search, going to a balanced search and different covariance matrices for each group. Stacking therefore seems to be preferable to the singles. By comparing three different balanced

search settings with a pooled covariance matrix, unbalanced, and unbalanced with a separate covariance matrix for each group, we can see that most classifiers improve their performance with an unbalanced search, while, as was rather predictable, for QDA the best performance is achieved by considering the matrices separately for each group and TRE also achieves its best performance.
Even if there are no outliers it is interesting to see the posteriors and cross validation probabilities plots because they are useful tools for data exploration. In fact, each classifier displays different behaviour, and in its function of exploring data as well as outlier detection, the Forward Search is very powerful. There do not seem to be any outliers but obviously there are some units that enter the last steps of the search which are less in agreement with the model, in particular unit 26, which is the last of all to join the subset. Basically the allocation probability for the two groups is similar, and there are better classifiers, but they are all quite good in this situation. LDA and QDA have quite similar levels of correct allocation. See Figure B1-B2 in Appendix to Chapter 6.

Generally, Forward Plots of the posterior probabilities of belonging to the correct population may display either confused or clearly separate trajectories. If the classes are clearly separate in the posterior graphs, the trajectories which depart from the value of 1 should not be either many or confused, but they should all be near to 1. If the populations overlap, and thus correct classification is difficult, there will also be confusion in the trajectories because they represent the posterior probabilities of each unit with respect to the relative population. They should all be near to 1 for the first grop and near to 1 for the second group, apart from the outliers which will enter at the end of the search.

6.3 Simulated contaminated data

The second dataset on which the Stacking procedure is applied in a Forward Search context has the same characteristics as the one used in the previous section but it has undergone a contamination of the first twenty observations belonging to group 1 through the addition of a constant c (c=+5) to the original value of each one of them.
On examining the scatterplot matrix in Figure 6.3.1 we can see that the situation appears to be very interesting, since the contamination of the group of observations is quite high compared to the degree of separation originally imposed between the two populations (δ=1.5), and the outliers of the first group are positioned to the right of the second group for all the pair wise scatter plots of the variables.

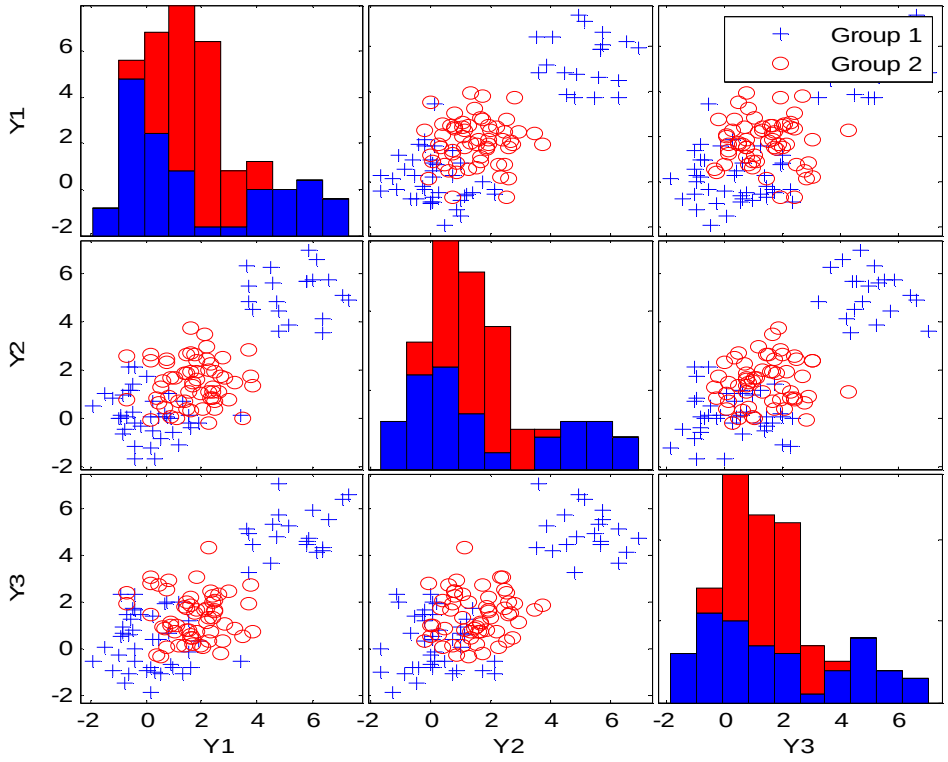

Figure 6.3.1 Simulated contaminated data. Scatterplot matrix with bivariate scatters of the three variables and histograms on the main diagonal. The units in Group 1 are represented by blue crosses.

As far as the rest of the observations are concerned, no group is perfectly separated from the others in any bivariate scatter.
Outliers and influential observations can be detected graphically and it is extremely useful to monitor the maximum Mahalanobis (see Figure B3 in Appendix to Chapter 6) distance for units in the subset and the minimum Mahalanobis distance among the units not belonging to the subset.
An analysis of the plots of the minimum and maximum Mahalanobis distances clearly reveals the entrance of anomalous observations into the subset, which is shown by a peak in the curve (in blue) relative to group 1, respectively, in the step prior to the inclusion of the first anomalous unit (for the minimum), with a sharp increase when the first atypical unit joins the subset (in the maximum). Therefore, since we use a balanced search, from step 80 anomalous values will begin to enter group 1. In particular, the first unit to enter is number 26, which is not one of the twenty observations on which we carried out the contamination, but in any case it is not in agreement with the units contained in the subset. Instead during the following steps, all the units that we contaminated enter group 1, and during an interchange at step 100 when two new units (1 and 11) enter the subset, unit 26 definitely comes out, and re-enters at step 116, thus confirming its distance from the “good” units of the subset with small Mahalanobis distances.

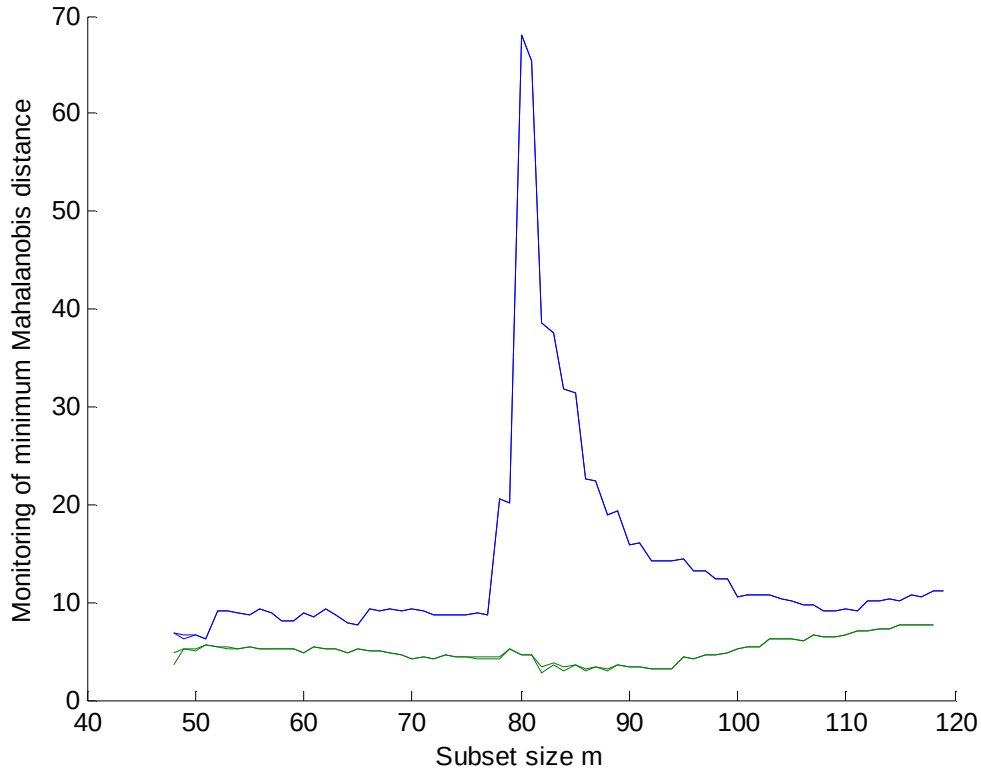

Figure 6.3.2. Simulated contaminated data. Plot of minimum Mahalanobis distances for units not belonging to the subset. Blu line, Group 1. Balanced search.

We are therefore interested in monitoring the effects of the entrance of outliers and of unit 26 into group 1 on the decision rule both of base classifiers and of Stacking. We start by analysing the cross validation error graph for the units inside the subset, above all at the point where we have established that anomalous values enter group 1 and then we analyse the effects of the classifer probabilities in the graphs.

In the cross validation error graph, in Figure 6.3.3., at the point where we have established that anomalous observations make their entrance, by means of the Mahalanobis distance plot, there is a general increase in the curves relative to classifiers but there is clearly an increase for those relative to Linear Discriminant Analysis and Logistic Regression. This particular condition of the two classifiers is quite interesting because we had already been able to verify an increase in cross validation error, especially for these two classifiers compared to the others, also in the traditional approach (in Section 4.3) in the presence of quite consistent contamination and only for one group.

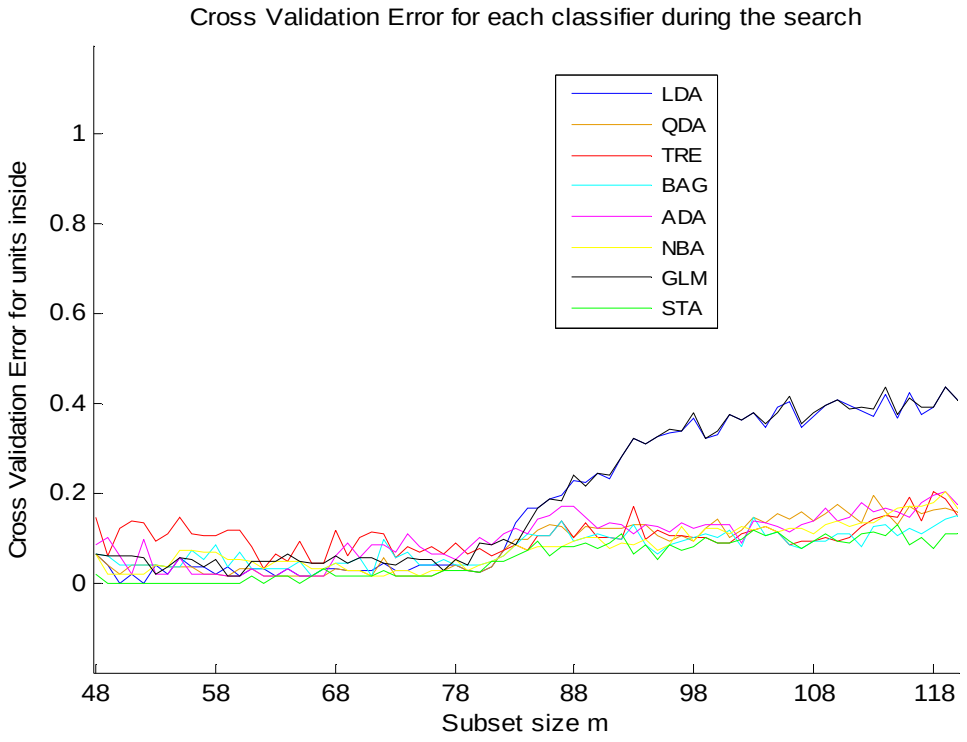

Figure 6.3.3. Simulated contaminated data. Cross-validation error for units belonging to the subset of the seven base classifiers and Stacking. Balanced search.

Moving on to the hypothesis of contaminating two classes, as we had already seen in Section 4.3, the four parametric classifiers in any case achieve a high level of error, as can be seen in Appendix to Chapter 6, Figure B4. It should also be remembered that this analysis was carried out using the default setting already described in the previous example, that is to say a balanced search and a common covariance matrix for both the groups, which represents the basic condition for LDA. We will subsequently see what the effects of modifying such choices will be. The graphs both of the posterior probabilities and cross validation probabilities are very interesting, since they show that all the classifiers are sensitive to the entrance of the outliers, but above all that, all of them respond in different ways and there are changes of direction. The behaviour of unit 26 is always interesting (already seen in the previous example) as it is always misclassified until the end of the search, thus confirming what has been said previously. By analysing the plot of the posterior probabilities of Linear Discriminant Analysis, we can observe there are not units that are perfectly classified until the end of the search. When the outliers enter, the trajectories undergo a consistent change of direction, and they are concentrated on the last steps around 0.5.

By comparing the graph of LDA posterior probabilities on contaminated data in Figure 6.3.4 with the corresponding plot on non-contaminated in Figure B1 in Appendix to Chapter 6, it is clear that there is a considerable effect caused by the outliers on this classifier in terms of performances.

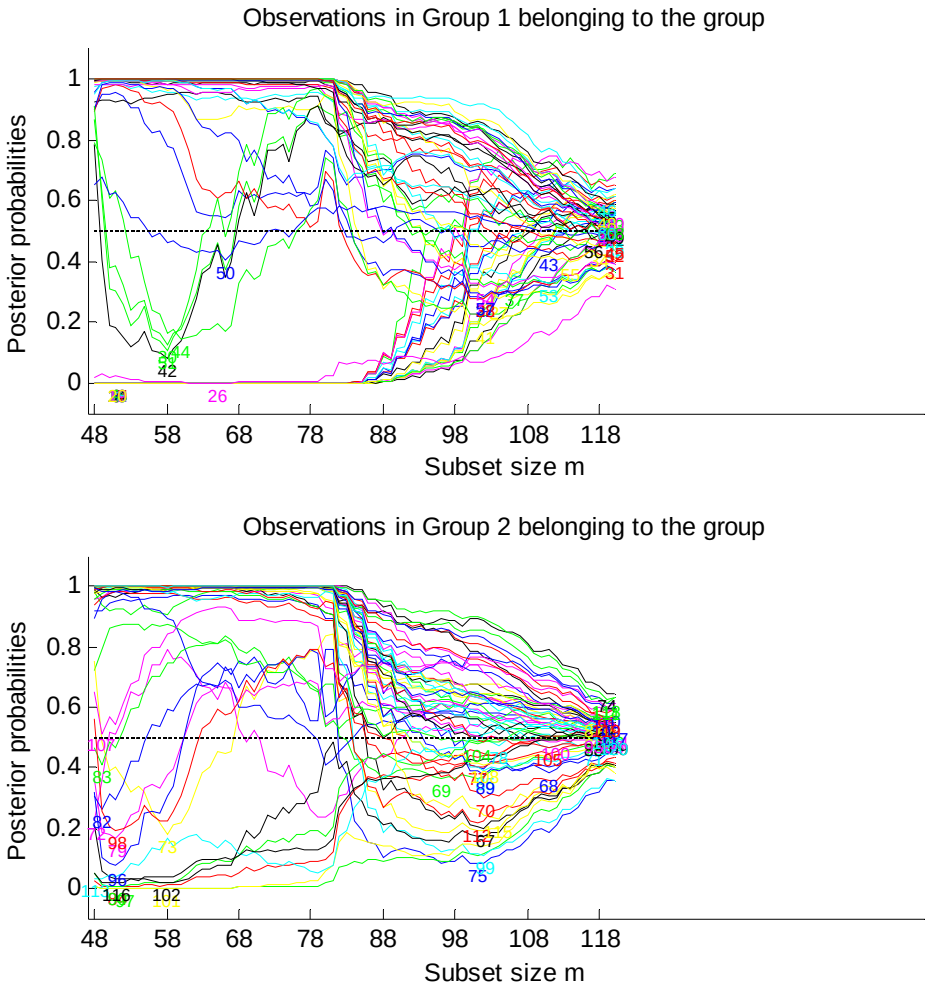

Figure 6.3.4. Simulated contaminated data. Linear Discriminant Analysis: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.

This is perfectly in line with what we have seen in the plot of cross validation error for units inside the subset that displayed quite a consistent distortion of the allocation rule. The decision rule, therefore, is distorted, because of the entrance of outliers, and the same thing happens in the plot relative to cross validation probabilities, which in fact are sometimes worse. Since the posterior probabilities also concern units that take part in the fit of the model: the model is fitted better compared to what was forecast. The units that take part in the fit are the best ones, that is to say those which form the central part of the distribution. The behaviour of Quadratic Discriminant Analysis is different and decidedly better, and despite the presence of outliers, is able to correctly classify some units for each step of the search. However, even for this classifier we can see the effects of the entrance of units that are not in agreement with the model, which are quite clear if we compare them with the corresponding plot that refers to the non-contaminated dataset; from step 80 there is a change of direction for a consistent number of trajectories of Group 1, which has the effect of worsening the performances of the classifier by

causing an increase in the number of units that are either completely misclassified or not perfectly classified at the end of the search. In the second group, from step 80 there is a change in the trajectories of the units and there will be no more perfectly classified observations, that is to say with values close to 1. Furthermore, a considerable number of units will be misclassified for most of the search and will only be attributed to Group 2 in the last steps, while others will be completely misclassified.

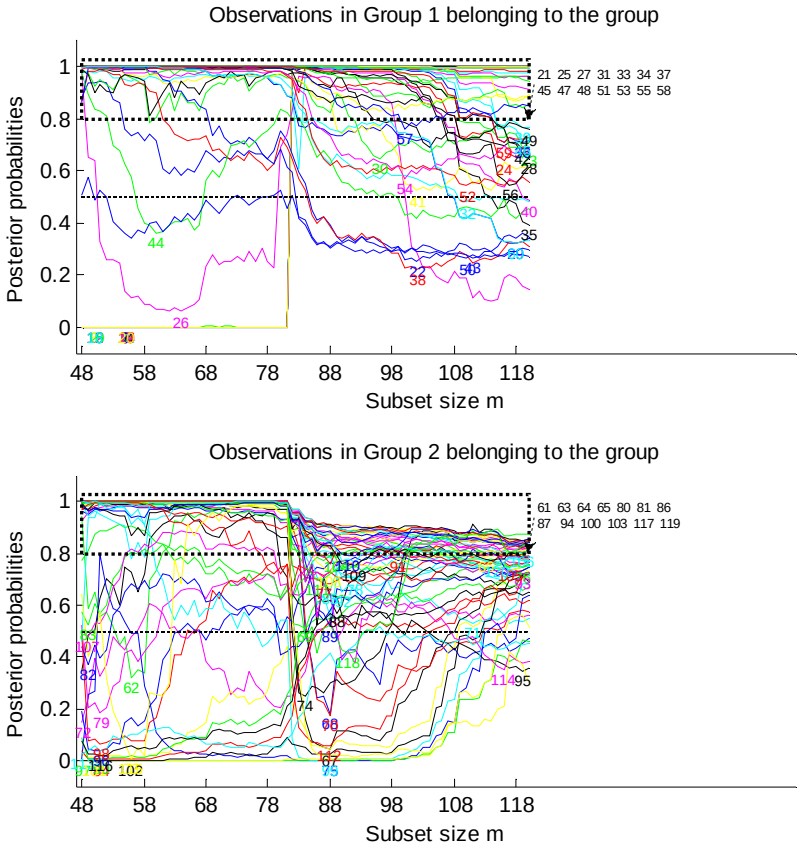

Figure 6.3.5. Simulated contaminated data. Quadratic Discriminant Analysis: posterior probabilities of correct classification of the units in Group 1 (upper panel) and in Group 2 (bottom panel). Balanced search.

A filter has been created for the posterior and cross validation probabilities graphs. It is used to filter the trajectories that were under the threshold of 0.8 at least for 1% of the number of total steps. For these trajectories we will show, in correspondence with the minimum point reached, the indication of the unit referred to. The non-filtered trajectories, or rather those that will stay over 0.8 for 99% of the steps, are shown by a box with dotted lines. At the side there is a caption which reports the units that are always correctly classified and inserted in the box.
Interesting behaviour is displayed by the Classification Tree, which is generally considered to be a “weak” learner. It gives better performances than the two parametric methods we have just seen; it is better able to classify the units of the second group, and it achieves a discreet performance for those of the first group.

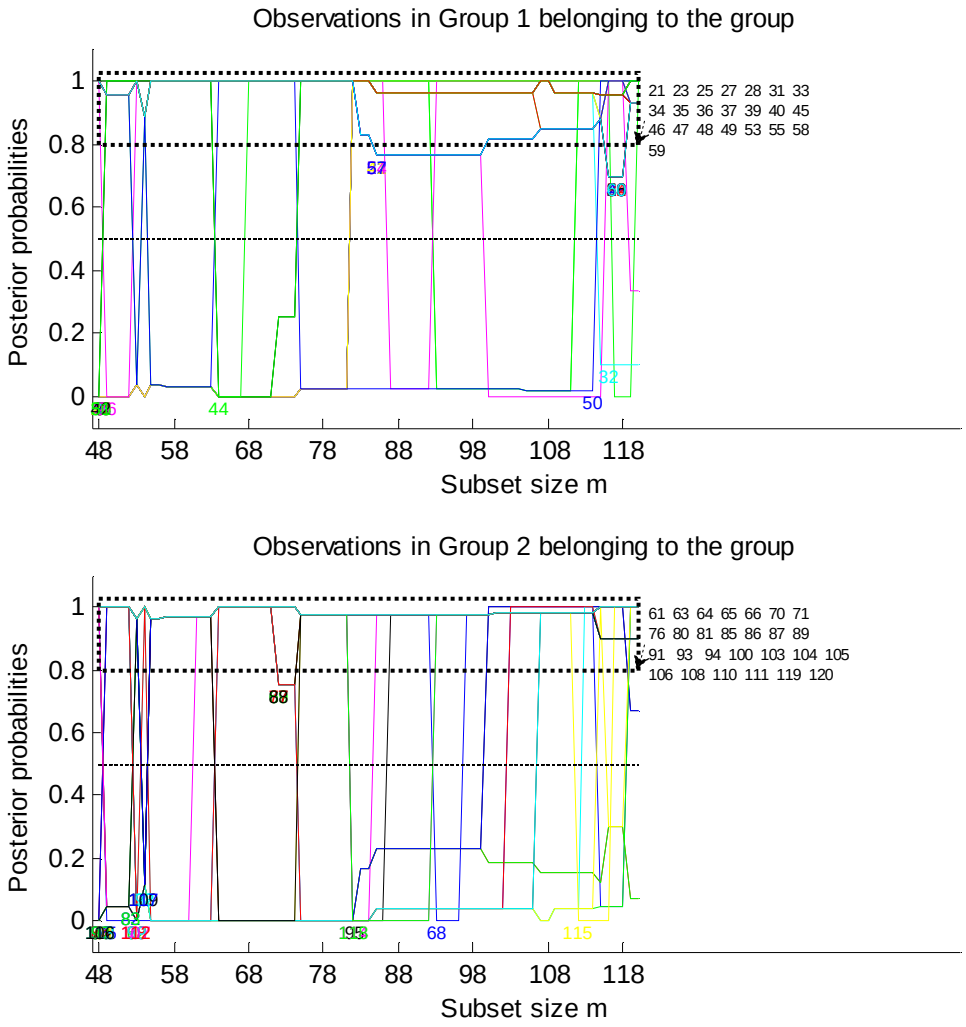

Figure 6.3.6. Simulated contaminated data. Classification Tree: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

The behaviour of the Bagged Classification Tree in Figure 6.3.7 is very interesting; it gives a good performance, as it is able to correctly classify a considerable number of units.
The performance is better, of course, in the second group compared to the first (in which the outliers enter directly): however, towards the end of the search it is able to classify almost all the units in the first group with the exception of one. The behaviour of unit 26 should also be pointed out, which in this case is mostly allocated to Group 1 (it is the first unit of group 1, whose entrance, at step 80, is indicated by a peak in the plot of the maximum Mahalanobis distances) and in the last steps, which correspond to its re-entry into the subset (at step 116, having presumably gone out at step 100), it is finally assigned to that group.

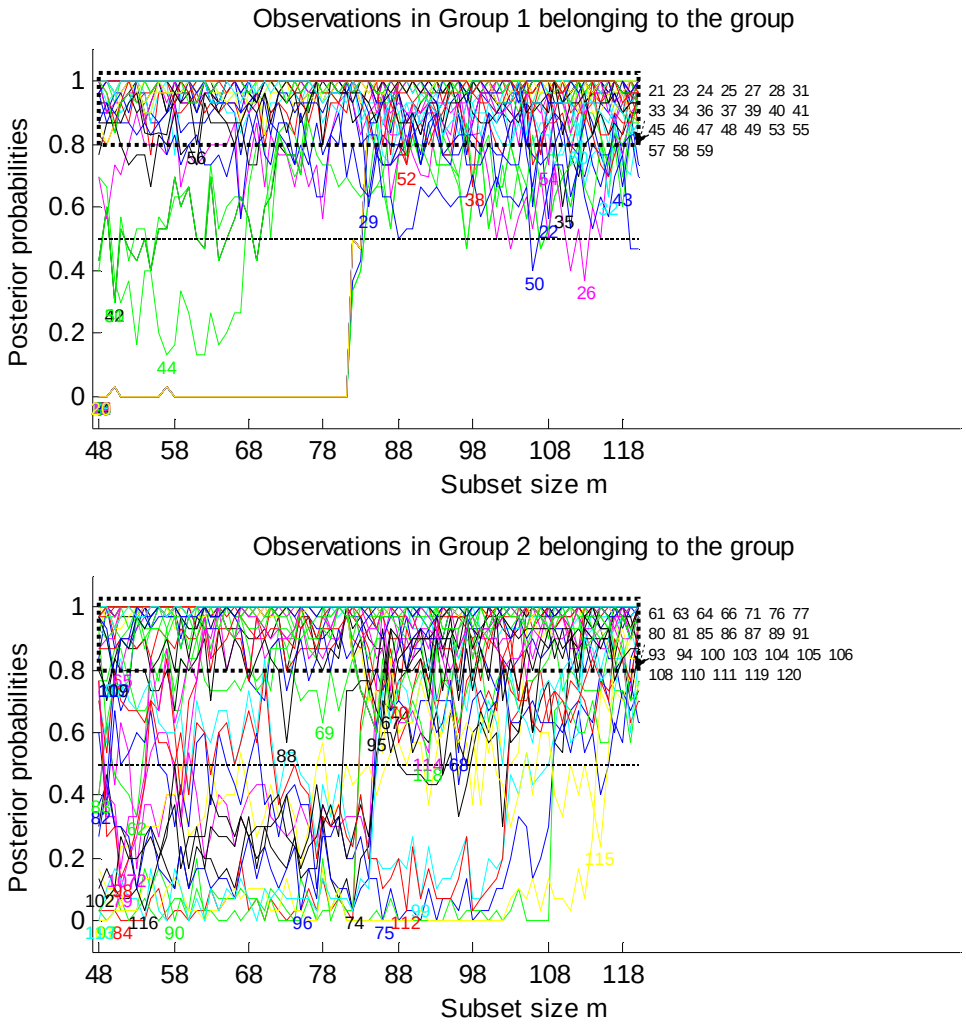

Figure 6.3.7. Simulated contaminated data. Bagged Classification Tree: posterior probabilities of correct classification of the units for Group 1 (upper panel) and Group 2 (bottom panel). Balanced search

The behaviour of the Adaboost classifier in Figure 6.3.8 is particularly interesting. Before the entrance of anomalous values, the allocation probabilities seem quite high for both groups. With the entrance of the first value that is not in agreement with the model (unit 26), the trajectories of the units of the first group show a sharp decline and they all settle around 0.5 at subsequent steps of the search until the entrance of the last anomalous value. The presence of outliers therefore generates a situation in which there will never be any perfectly classified units, but they will mostly be allocated to group 1 anyway, with moderate probability. The trajectories relative to the units of Group 2, which initially seem to be well classified, undergo an even sharper decline than in the other group, but they will all settle above 0.5 and will therefore be allocated to class 2 (including unit 115 which enters at step 117 and therefore has quite a big distance), with the exception of unit 101 which makes its entrance at step 115.

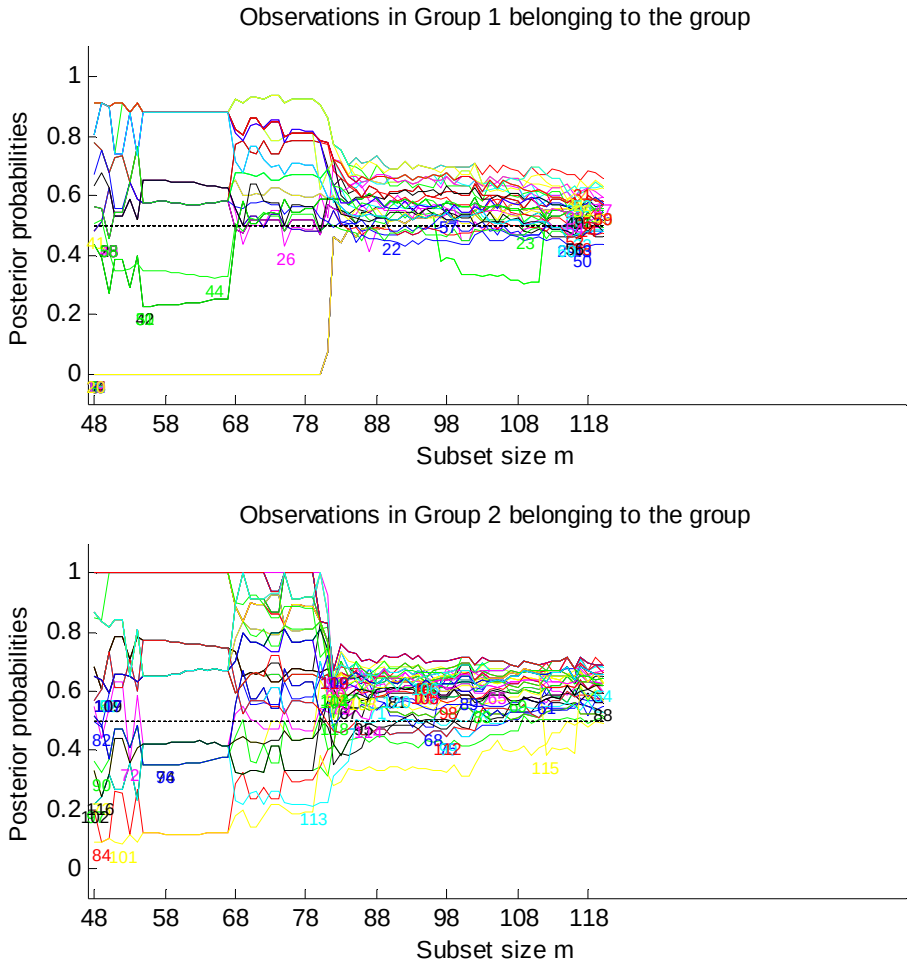

Figure 6.3.8. Simulated contaminated data. Adaboost: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

By comparing Figure 6.38 with the one relative to the cross validation probabilities given by Adaboost in Figure 6.3.9, it can be seen that with the entrance of the outliers in the prediction phase, the units from Group 1 that are allocated to it are much fewer, and most of the units will be misclassified. Thus, the model is fitted better than it is able to predict (as was predictable, since in the posterior probabilities there are also those that refer to the units that have taken part in the fit of the model), the classification rule will be distorted also in the case of the cross validation probabilities. For the second group, for which many units were correctly classified in the first steps of the search, the number of misclassified units at the end of the search is decidedly lower compared to the first group, but higher when compared to the situation shown by the plot of the posterior probabilities.

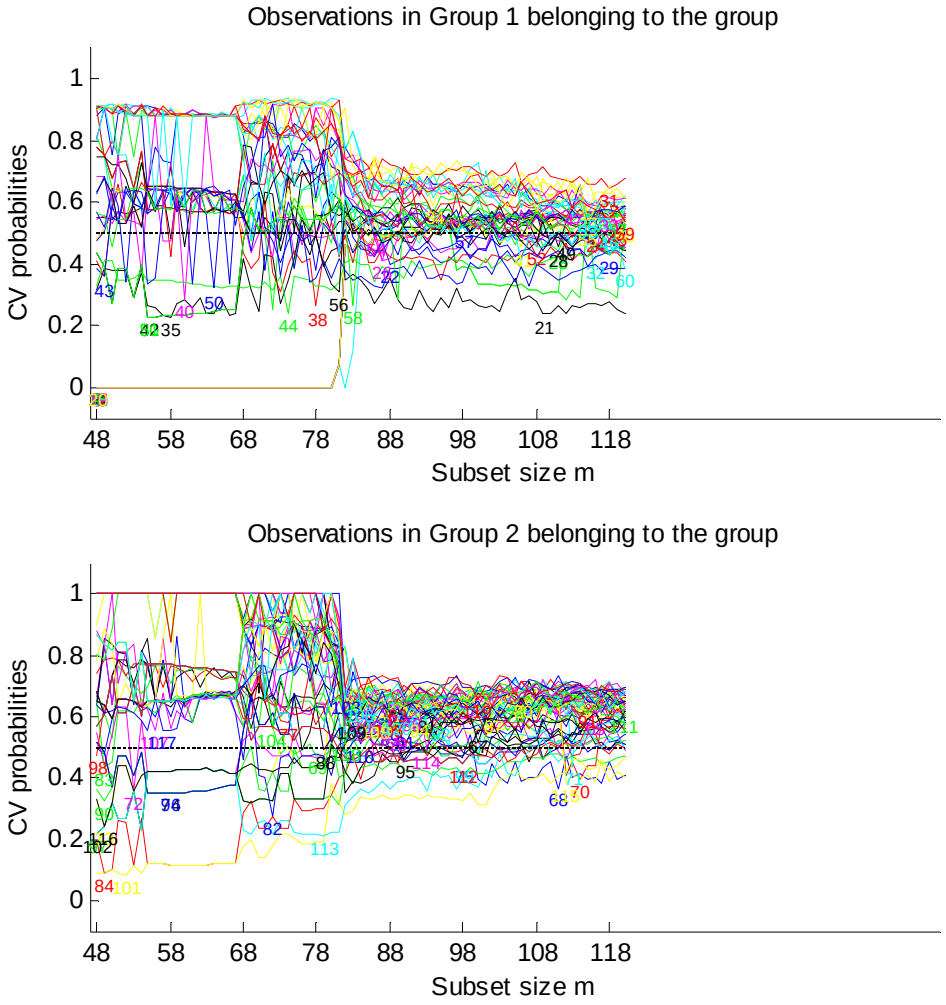

Figure 6.3.9. Simulated contaminated data. Adaboost: Cross-validation posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

When observing the plot of the posterior probabilities given by Naive Bayes, the effect caused by the entrance of units not in agreement with the model is very clear. With respect to a number of units corretly classified both in the first steps and during the whole search, in Group1, when the outliers enter, there is a clear change of direction in the trajectories and at the last step a lot of units are misclassified, and for the others allocated to the group, the probability of allocation diminishes dramatically. In Group 2 the algorithm is able to better classify the units compared to Group 1, and the number of units allocated with a probability of ≥ 0.8 is actually higher, even if there are units that are perfectly classified only as far as step 80. In the last steps of the search, all the units except three, 101, 115 and 75 (which enter, respectively, into steps 115, 117 and 119) will be allocated to Group 2.

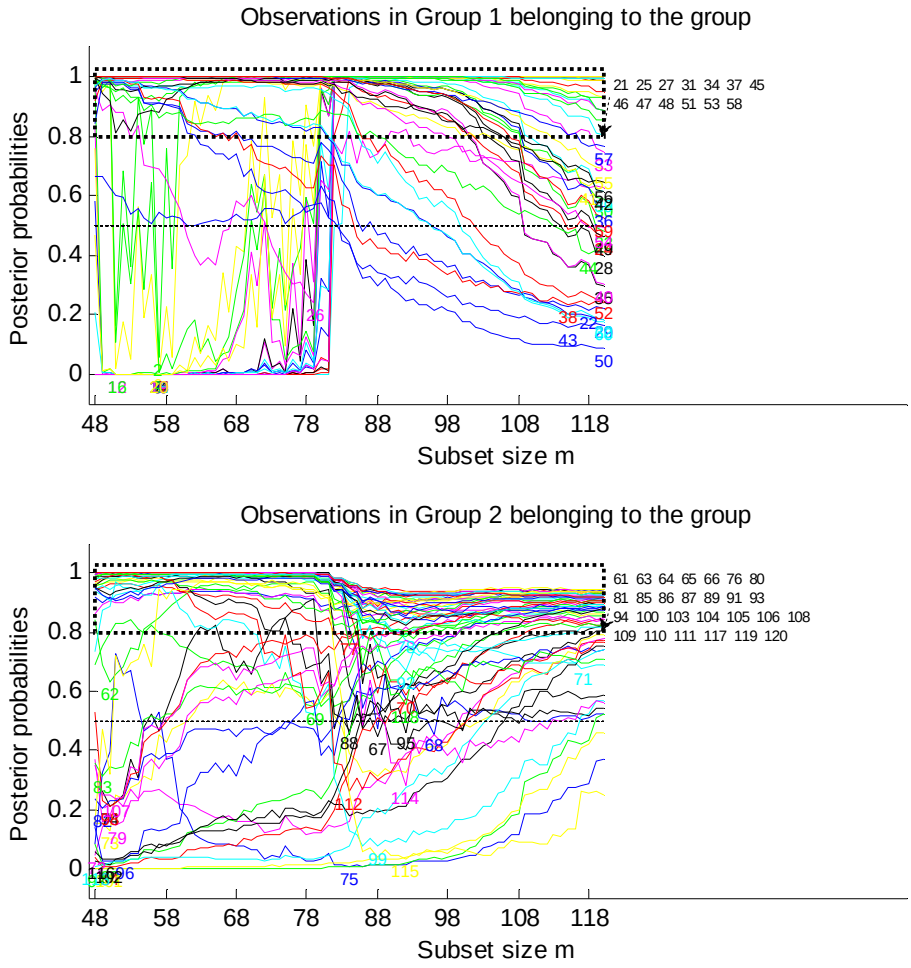

Figure 6.3.10. Simulated contaminated data. Naïve Bayes: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

The behaviour of Logistical Regression seems interesting, as we can see from the plot of the cross validation error, and together with LDA it is one of the classifiers with the worst performance in the presence of outliers. The entrance of outliers can clearly be seen from the plot of the posterior probabilities in Figure 6.3.11. there is a clear degeneration in the probabilities and, after the entrance of the outliers, no unit in either group will be perfectly classified and the proportion of units that will be misclassified is quite high.
The situation is also very similar regarding the cross validation probabilities which we therefore omit for the sake of brevity.

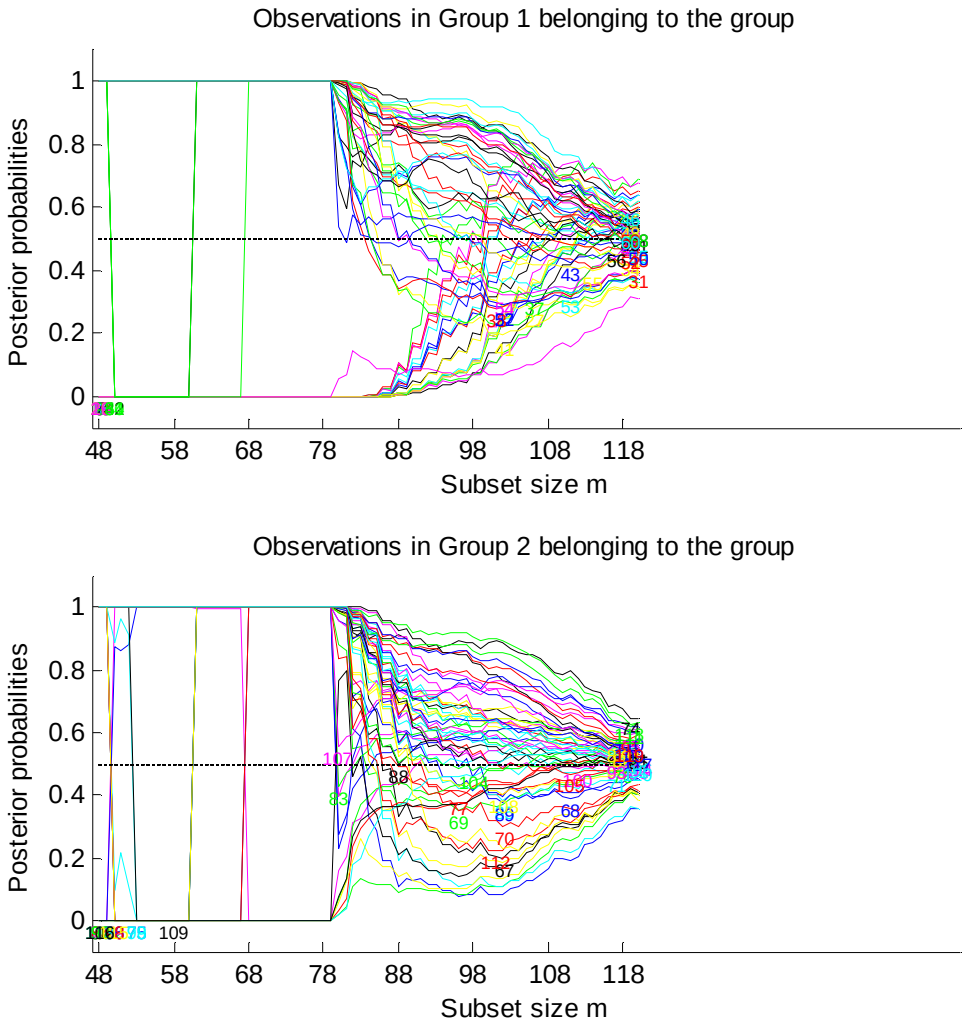

Figure 6.3.11. Simulated contaminated data. Logistic Regression: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.

Stacking is confirmed as the best classifier as can be seen in Figure 6.3.12, since it is able to correctly classify the highest number of units compared to the other classifiers and in the last steps of the search all the units will be allocated to the respective class. We can also appreciate this classifier, because of the effect it has on the entrance of the outliers on the classification rule. Of course it classifies the units of the second group better.

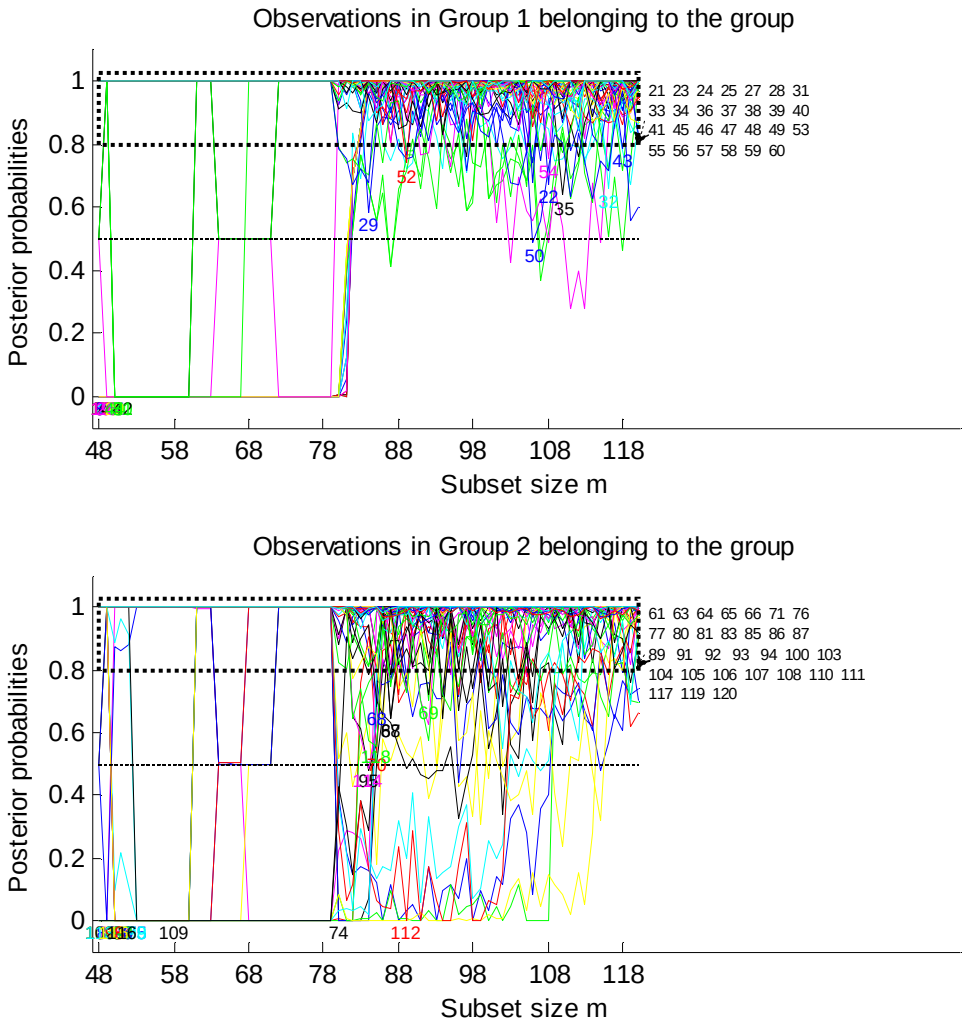

Figure 6.3.12. Simulated contaminated data. Stacking final classifier: posterior probabilities of correct classifcation of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.

By starting from an observation of the plots of Mahalanobis distances, which show the entrance of the outliers graphically, we can evaluate the alteration of the performances of the classifiers, or rather the distortion of the classification rule by verifying at the entrance point the effect caused by atypical observations of the plot of cross-validation errors for units inside the subset and in the plot of the posterior and cross-validation probabilities that show the probabilities for each unit of belonging to the relative group at every step of the search.

In the example we have just seen, the effect of the outliers present in Group 1 has been quite consistent if compared with the situation in which the dataset was not contaminated, especially for LDA and GLM, which have a very high prediction error. In this simulated example, since

the contamination has been induced, we knew which units were altered, apart from unit 26 which is evidently quite distant and in most cases is misclassified. In the examples of application of the method to real data that we will see in the following sections, for which we have no a priori knowledge of the existence or the presence of the outliers apart from the graphic evidence provided by the forward plots of Mahalanobis distances, we will use a useful tool for the automatic detection of the outliers, even if this done separately for each population. It does not seem necessary in this case.

6.3.1 Simulated contaminated data in the hypothesis of a non-balanced search and a specific covariance matrix for each group

Since both balanced and unbalanced searches provide suitable information about the structure of the data, instead using a balanced search and a pooled covariance matrix will make it very interesting to explore different settings, as we will also see in the hypothesis of a non-balanced search and a specific covariance matrix for each group. Since we progressively include units with small Mahalanobis distances, in an unconstrained search the units of a group with small variance will tend to be included before those from a group with a larger variance. So, if a group contains some outliers, these will enter in the last steps of the search.
The effect produced by using different settings in the plot of minimum Mahalanobis distances, in Figure 6.3.1.1, is very interesting. There is an interruption in the curve of the minimum relative to Group 2. When looking at the units included in the subset at each step of the search (bearing in mind that at every step the new subset is compared with the old one and we can only see the units present in the new subset but not those in the old one), we see that up to step 99 the units of Group 2 (green line) have all entered, and for this reason there will be an interruption in the curve of the minimum relative to that group. Since, however, at each step there is the computation and ordering of Mahalanobis distances, at step 100 only units from Group 1 will enter (26 will be the first unit to enter, followed by the contaminated units), and one new unit at each step, up to step 107.
A masking effect will be produced, as the units of Group 1 all have small Mahalanobis distances, therefore from step 108 an interchange occurs and two or more new units of group one (outliers) enter the subset at each step, and units of group 2, which have entered into the last steps before the interruption of the minimum curve, then have to leave the subset in order for the size to increase by one unit.
Thus, when almost all the outliers have entered, the centroid of Group 1 will move so much that there will be a masking effect and the units in Group 2 will have greater Mahalanobis distances than those in Group 1, and from step 108 to step 113 they will come out to retrace their last steps together with unit 26 from the first group. By using a balanced search, we would never have been able to achieve this particular effect.

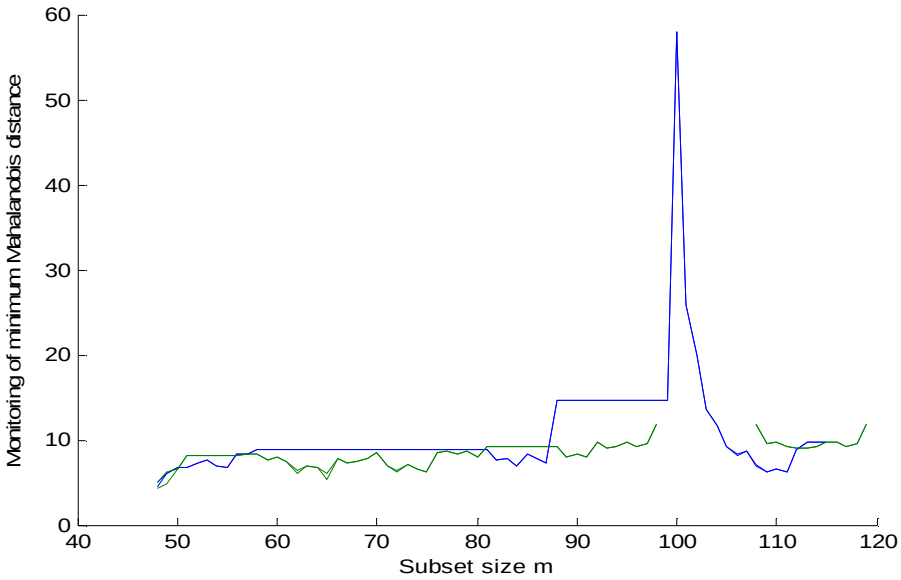

Figure 6.3.1.1. Simulated contaminated data. Plot of Mimumum Mahalanobis Distances for units not in the subset. Unbalanced search. Blue line, group one.

Figure 6.3.1.2 shows the increase in cross validation error, especially for the LDA and GLM classifiers when the outliers enter, which happens in the last steps of the search, since we have used an unconstrained search.

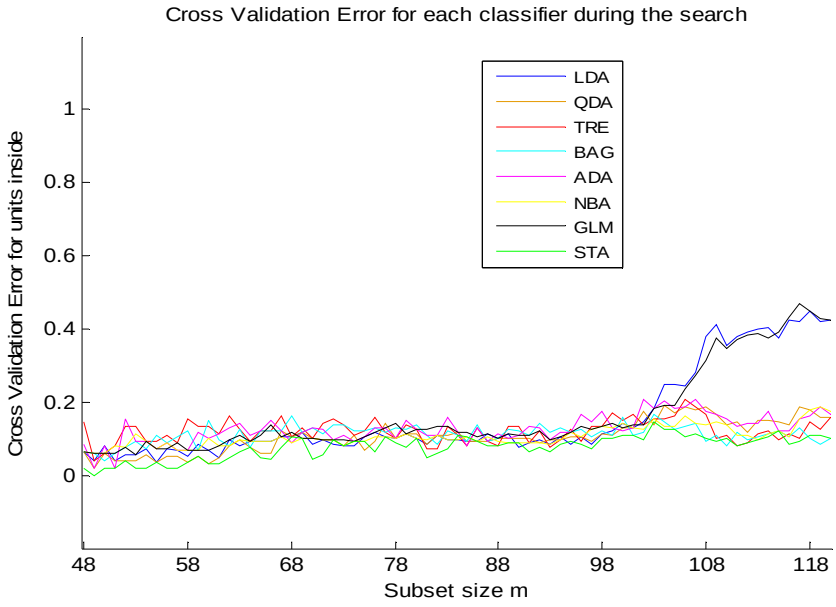

Figure 6.3.1.2 . Simulated contaminated data. Cross-validation error for units belonging to the subset of the seven base classifiers and Stacking. Unbalanced search.

The effect of the outliers will also be appreciated in the plots of the probabilities of allocation to the respective group relative to the single classifiers and to Stacking (which we omit here as it is very similar to those previously shown), which will display in general a decline in correspondence to the entrance of atypical observations, with the exception of Classification Tree which does not seem to be much affected by this and in the plot of the posterior probabilities is able to correctly classify a higher number of units in both groups compared to the other classifiers.

The behaviour of Adaboost is strange in this case too, since in the presence of outliers it is able to allocate the units to their respective class (completely in the second, less so in the first), but with much lower probabilities, only slightly higher than 0.5, thus demonstrating that in this way it is influenced by the anomalous units.

In the following sections we will see the application of the proposed Stacking scheme to real data in a Forward Search context.

6.4 Application to real data

6.4.1 Electrodes Data

The first application to real data is carried out on "Electrodes Data" (in Atkinson, Riani and Cerioli (2004)), already introduced in the traditional part. This is a dataset made up of five variables which represent measurements of electrodes produced by two machines which therefore represent two groups. Let's summarize the main characteristics:

$n_1 = 50$

$n_2 = 50$

$y_{1,} y_{2,} y_5 = diameters$

$y_{3,} y_4 = lengths$

The scatterplot matrix in Figure 6.4.1.1 shows a certain overlap between the two groups, especially for some variables, while in other cases the separation seems to be more substantial, for example for the variable y4. There also seem to be some anomalous values present.

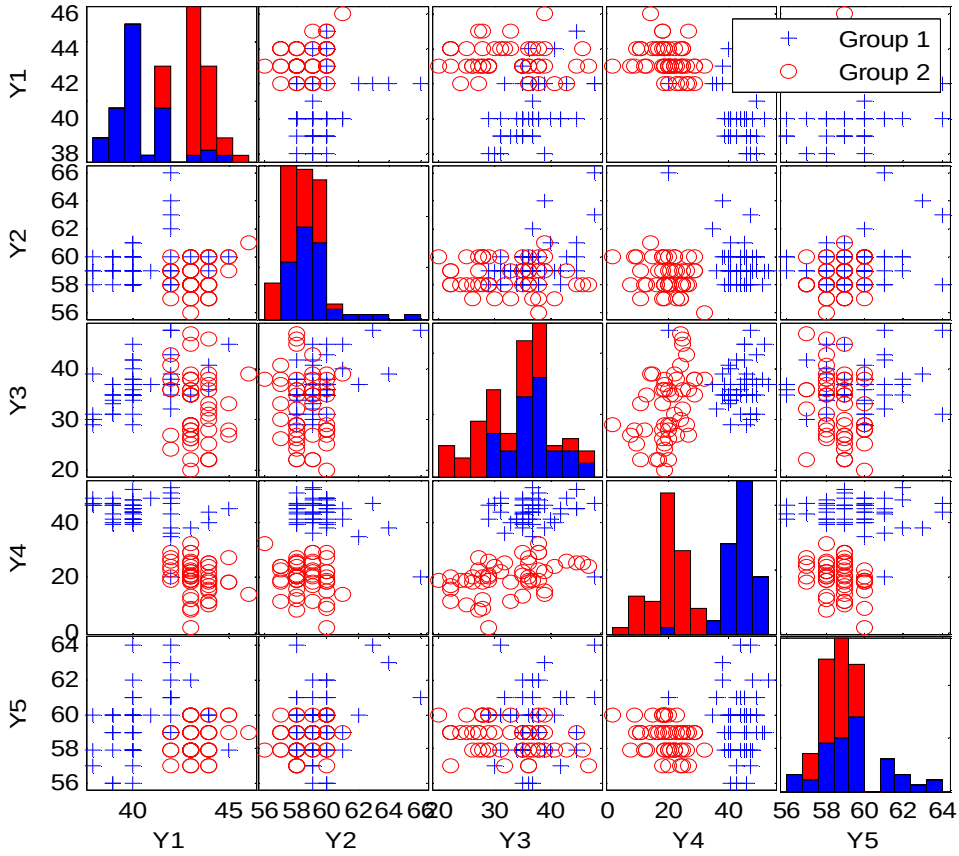

Figure 6.4.1.1 Electrodes data. Scatterplot matrix with bivariate scatters of the three variables and histograms on the main diagonal. Units in Group 1 are represented by blue crosses.

We are interested in exploring data and investigating the behaviour of the single classifiers and of Stacking on this dataset, in terms of differences in their performances and finally we will compare the graphs of the posterior probabilities and cross validation probabilities in order to see if there is a difference or an agreement between the fit of the model and prediction.

We start with a subset size of 40 out of the total 100 units, and a balanced (constrained) search and a common covariance matrix is used for the two groups. The subset will contain an equal number of observations from each group and the search will usually proceed with the insertion of one unit for each group.

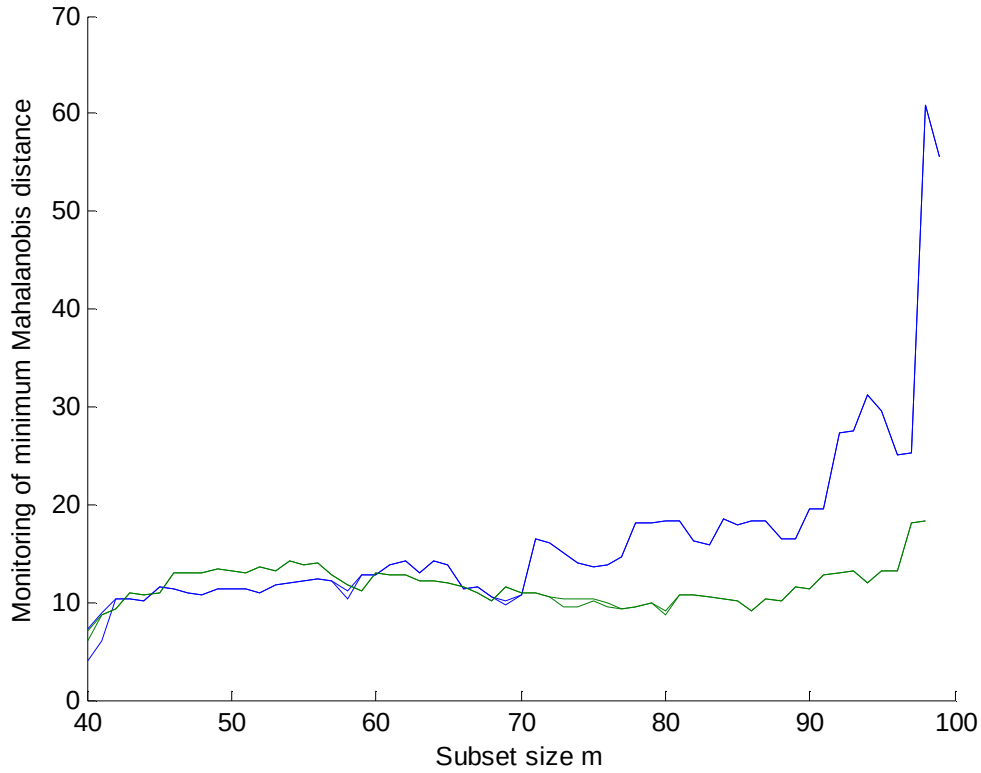

Figure 6.4.1.2. Electrodes data. Plot of minimum Mahalanobis distances for units not belonging to the subset. Blu line, Group 1. Balanced search.

The plot of the minimum Mahalanobis distance of the units outside the subset in Figure 6.4.1.2 reveals the entrance of anomalous observations into the subset which is shown by an increase in the curve (in blue) relative to group 1 in the step prior to the inclusion of the first anomalous unit, $m = 75$ and it seems that there is a larger distance for units of this group in the second part of the search. We will use also an unbalanced search subsequently.
In a non-balanced search the observations with smaller Mahalanobis distances join earlier the subset. From the plot of minimum Mahalanobis distances (see Appendix to Chapter 6), referred to the unconstrained search, we can see that from $m = 96$ only units of Group 1 enter the subset. Some units from Group 1 (with larger Mahalanobis distances than the other ones) will be, therefore, misclassified until the last steps of the search and they will be allocated to Group 1 later compared to a balanced search.

The cross validation error plot in Figure 6.4.1.3 shows that the performances of the classifiers are generally very good, the error is generally quite low and the highest values are achieved by Adaboost, Bagging, Classification Tree and by GLM and NBA.

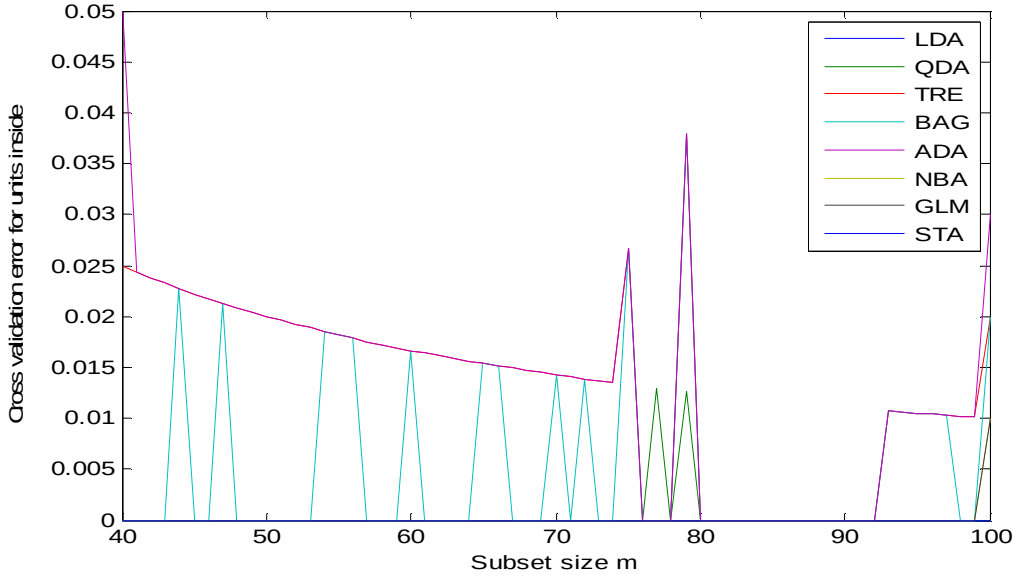

Figure 6.4.1.3. Electrodes data. Cross-validation error for units belonging to the subset of the seven base classifiers and Stacking. Balanced search.

In terms of full allocation of the units to the respective group (with probabilities ≥ 0.8 that is therefore distinguished from the allocation to the respective group for which a probability > 0.5 is sufficient) we can say that as far as a balanced search is concerned there are 3 classifiers in which the posterior probabilities coincide with the cross validation probabilities (Quadratic Discriminant Analysis, Naive Bayes and Stacking). See the plots in Appendix to Chapter 6, Figure B5-B10. In the unbalanced search, however, there will be five classifiers (Linear Discriminant Analysis, Quadratic Discriminant Analysis, Naive Bayes, Logistic Regression and Stacking) for which this is verified as can be seen in Appendix to Chapter 6, Figure B11-B20. Therefore, in terms of full allocation, there is a coincidence between fit and prediction. In both cases, however, the units of the second group are better classified.
The units that appear in the caption for the graph refer to the full allocation, but of course for the assessment of the performance of a classifier we will also take into consideration the allocation of the units to a group also with a probability of only > 0.5, above all in the hypothesis that the groups are not (as in this example) perfectly separable for all the pairs of variables. As we have already indicated, if the groups were perfectly separated the trajectories should all be near to 1. The performances of the classifiers are quite similar as we had imagined from observing the plot of the cross validation error for the units inside the subset, which shows a very moderate error equal to 0 at every step of the search for LDA and Stacking (which is not one of the best for perfectly classified units) and for the almost the sum total of the search for QDA; from the analysis of the plot of the posterior probabilities (which we omit to report for the sake of brevity and which are in the Appendix to Chapter 6) for almost all the classifiers the sum total of the units is allocated to the relative group, even if some units (which we can consider as anomalous observations) are allocated to the respective class in an advanced stage of the search (later in the case of the unbalanced search). Particular attention should be paid to unit 9 (the last unit to enter the subset) which will always be misclassified until the last step of the search when it will be allocated to Group 1, for all the classifiers, except for Adaboost where it

will remain completely misclassified. Furthermore, for Adaboost, some units for Group 1 and Group 2 will conclude the search by being allocated to the respective groups, but not completely, thus degrading the performance at the last step of the search and coming close to the behaviour analysed in the previous example (positioning a little higher than 0.5) which, instead, did not hold for all the search. By looking at the cross validation errors graph we can see that the highest error is achieved by Adaboost then by Bagged Classification Tree, Classification Tree and then Logistic Regression and Naive Bayes. The errors presumably derive from the non-correct allocation of the units to the respective groups by means of the cross validation probabilities. In this case the classifiers Classification Tree, Bagged Classification Tree and Adaboost (Figure B22, B24, B26 in Appendix), will not correctly allocate unit 9 (to Group 1) and unit 56 (to Group 2). For Adaboost, furthermore, in the cross validation probabilities there will also be a worsening of the allocation probabilities (Figure B6 in Appendix) for a group of units from Group 1 and Group 2, while Logistic regression and Naive Bayes will not allocate unit 9 to the first group.
Therefore, in terms of full allocation, in the balanced search there are no significant differences between posterior and cross validation with the exception of Adaboost and Bagged Classification Tree for which the model is in the prediction phase, able to fully allocate (probability ≥ 0.8) a number of units higher than that achieved by means of posterior probabilities, while in the unbalanced search, this is only the case for Adaboost. There is, however, a difference in terms of unit allocation (and therefore of the performance of the classifier) and in the posterior probabilities with the exception of Adaboost all the units are allocated to the relative group, even if belatedly, at the last step (as for unit 9 in the first group) and even if it is not done perfectly.

The situation is not the same for cross validation probabilities. Only LDA, QDA and Stacking will allocate all the units to the respective groups even though not perfectly (LDA) or in the last steps both in balanced and in unbalanced search.

In the following Section we describe briefly another two applications of the proposed Stacking scheme to real data in a Forward Search context for which a transformation of the data is necessary given the structure of the data.

6.4.2 Muscular Dystrophy Data and Off-the-book Workforce Data

In this Section we describe briefly another two applications of the proposed Stacking scheme to real data in a Forward Search context for which a transformation of the data is necessary given the structure of the data.

The first application to real data is carried out on "Muscular Dystrophy Data" (in Atkinson, Riani and Cerioli (2004)). We apply our stacking scheme implemented in a Forward Search context to complete dataset originally given by Andrews and Herzberg (1985) consisting of 194 observations divided in two groups (Group 1= non carriers, Group 2= carriers) and six variables.

The groups are not balanced:

n_1 =127

n_2 =67

The variable used are:

y_1 =age

y_2 =month of the year

y_3 = level of creatine kinase

n_1 =level of hemopexin

y_5 =level of lactate dehydrogenate

y_6 =level of pyruvate kinase

We start with a balanced search on untransformed data, and most of the classifiers achieve reasonable performances. However, the effect of the entrance of outliers is clearly visible in the plot of minimum Mahalanobis distances in Figure B27 in Appendix to Chapter 6 which shows quite a high distance between the two groups, as well as very high values for the curve relative to the second group (which suggests the use of an unbalanced search, but above all a timely transformation of data).

The effect of the entrance of outliers in the plot of posterior probabilities of Naive Bayes in Figure 6.4.2.1 is also quite clear. The behaviour of Stacking in Figure B28 in Appendix to Chapter 6 is interesting as it is able to allocate all the units to the second group. Furthermore, we used the automatic outlier detection procedure proposed in Riani, Atkinson and Cerioli (2009), even if this was done separately for each population that reveals the presence of outliers both in the first group and in the second one. Observations that appear such as outliers in untransformed dataset can be not outliers for some transformations of the data. Therefore, we decided to apply the method also to transformed data. We run a search using the suggested transformation (in Riani and Atkinson, 2001) with $\lambda_R = \left(-0.5, 1, -0.5, 1, 0, 0\right)^T$.

In the case of transformed data, there is an improvement in the plots of the posterior probabilities for all the classifiers and, therefore, an improvement of the decision rule. The detection of the outliers reveals the presence of a single outlier, unit 78 .

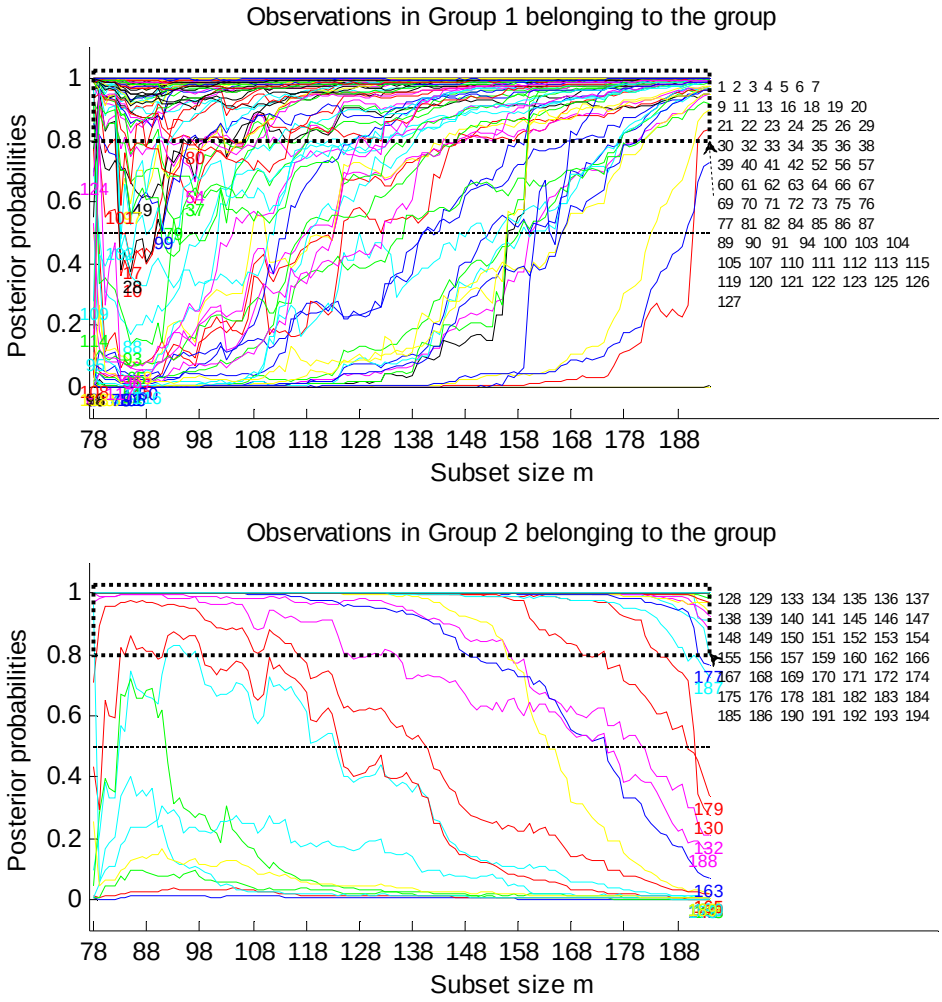

Figure 6.4.2.1. Muscular Dystrophy data. Naive Bayes: posterior probabilities of correct classifcation of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.Untransformed data.

The last application is carried out on the Off-the-book Workforce dataset already illustrated in the traditional approach. As we have already seen in Figure 4.61. in Chapter 4, also after the logarithmic transformation, we were dealing with a dataset with a marked level of overlap between the two groups and the presence of collinearity between variables. The plot of Mahalanobis distances in Figure B29 in Appendix to Chapter 6 reveals the potential presence of anomalous observations. The application of the Stacking scheme in a Forward search context confirms that this not a suitable dataset for classification because in all the plots of the posterior probabilities there is confusion in the trajectories and the classifers are unable to allocate the

observations to the respective groups. The only classifier that is able to classify some units both in a balanced and a non-balanced search and is able to assign all the units in the respective group in the last steps is Stacking in Figure 6.4.2.2.

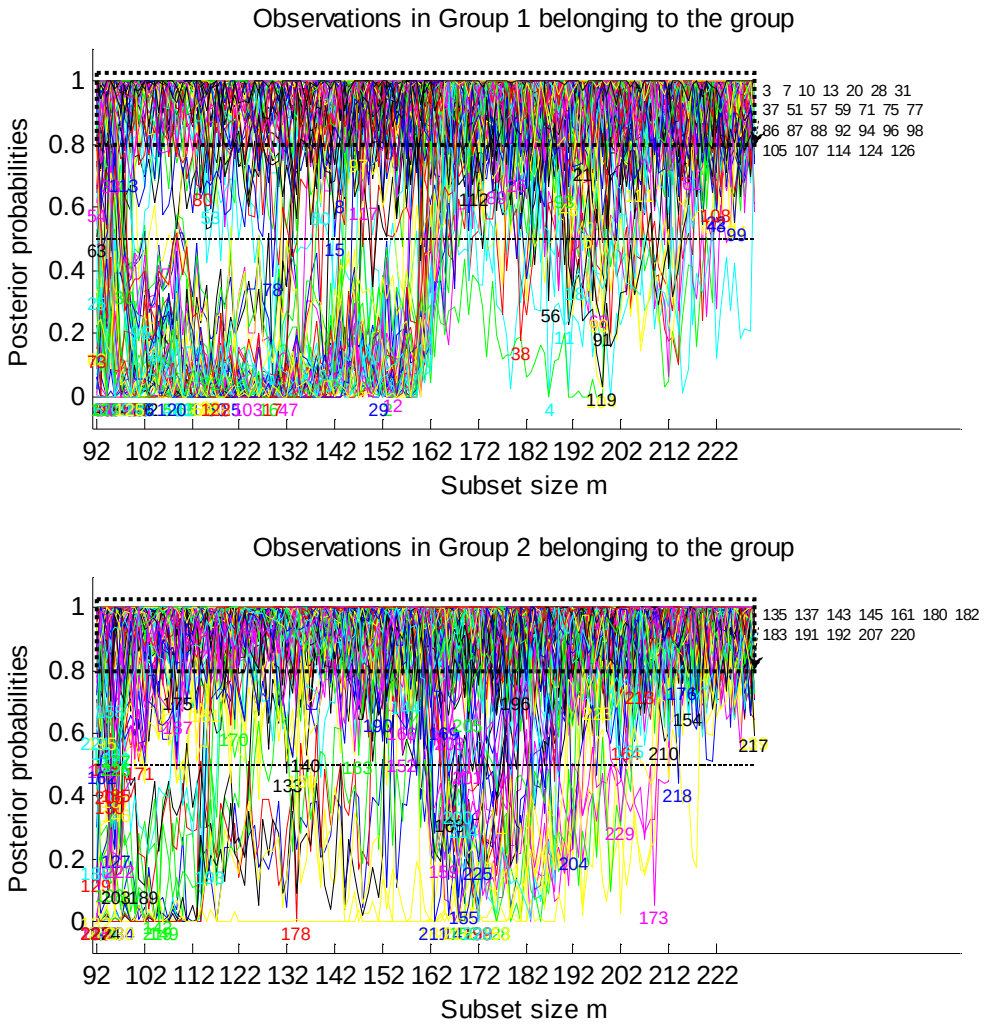

Figure 6.4.2.2. Off-the-book workforce data. Stacking scheme: posterior probabilities of correct classifcation of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.Untransformed data.

This is in line with what we have seen in the plot of cross validation error for units inside the subset in Figure B30 in Appendix to Chapter 6that displayed quite a consistent distortion of the allocation rule Stacking seems to be preferable in terms of cross validation error and is the best or among the best at every step of the search. The procedure for the automatic detection of outliers, even if this was done for each group separately, has revealed the presence of ten anomalous observations (units 4, 16, 38, 47, 56, 68,8 1, 90, 118, 123) belonging to group one

(absence of off-the-book workforce). Analogously with what we saw in the traditional part, this leads us to believe that there may be outliers that represent companies that have declared an absence of off-the-book work but give a signal of the potential presence of off-the-book employment in those companies.

Furthermore, for both of the applications for which we have no a priori knowledge of the presence of the outliers apart from the graphic evidence provided by the forward plots of Mahalanobis distances, we used the automatic outlier detection procedure proposed in Riani, Atkinson and Cerioli (2009), even if this was done separately for each population, since at the moment it is not possible to extend it to classification task. Since the main results obtained from the application of the proposed automatic outlier detection procedure, even if separately for each group, appear to be very interesting, the extension to problems of supervised classification could be very profitable.

6.5 Discussion

In the passage from the traditional approach to Forward Search, we acquire a much wider knowledge of the structure of the data, the trajectories of the observations, the importance exerted by each of these ones, whether outlier or not, on the model and on the aspects that are of interest for evaluating the performances of the base classifiers and of the Stacking scheme. With the traditional modality we could only see a static result, which is equivalent to the last step of the search but no information about the effect exerted by the single units, since they were treated simultaneously in the analysis, even if there was a logic of sample rotation such as cross-validation.

From the analysis of the results of some applications of the Stacking scheme proposed in a Forward Search context, referred to the simulated and real data, starting from observation of the plots of Mahalanobis distances, which show the entrance of the outliers (or influential observations) graphically, we can evaluate the alteration of the performances of the classifiers, or rather the distortion of the classification rule by verifying at the entrance point the effect caused by atypical observations in the plot of cross-validation errors for units inside the subset and in the plots of the posterior and cross-validation probabilities that show the probabilities for each unit of belonging to the relative group at every step of the search. It seems quite clear that the classifiers display different behaviour and that there is a considerable effect caused by the outliers on the classifier in terms of decline in the performances.

In the last two applications to real data, we used the automatic outlier detection procedure proposed in Riani, Atkinson and Cerioli (2009), even if this was done separately for each population, since at the moment it is not possible to extend it to classification task. Since the main results obtained from the application of the proposed automatic outlier detection procedure, even if separately for each group, appear to be very interesting, the extension to problems of supervised classification could be very profitable.

CHAPTER 7

Conclusions and Future Developments

In this thesis we were interested to investigate the predictive accuracy of one of the most popular learning schemes for the combination of supervised classification methods: the Stacking Technique proposed by Wolpert (1992), in particular, we made reference to the StackingC ensemble scheme as a starting point for our analysis, to which some modifications and extensions were made.
Since most of the research on ensembles of classifiers tends to demonstrate that this scheme can perform comparably to the best of the base classifiers as selected by cross-validation, if not better, this has motivated us to investigate empirically the performance of the single classifiers and of the Stacking technique, also in terms of stability and strength, by using two different approaches: one may be defined as traditional and the other as innovative.
We also focus on an aspect that has been covered much less by the studies and that in our view, however deserves special attention: the choice of the initial dataset. This is connected to the assumption that small changes in the dataset can lead to different models and that the presence of outliers might alter the parameters of the model. We were interested in investigating this aspect and in extending certain distinct elements in the Stacking scheme, as well as examining some characteristics neglected by the literature and also in proposing some distinct elements of the Stacking scheme. One of the main ones is the extension of the Stacking scheme in a Forward Search context which has enabled us to have a dynamic vision and monitor the behaviour of the classifiers and verify not only if one classifier achieves better performances than another one, but also if this is the case for all the steps of the search. This means being able to monitor the stability of the decision rule especially in the presence of anomalous values.
Starting from the recent advances proposed by the literature for the Stacking framework, the whole Stacking scheme was therefore implemented in Matlab and built in the form of an interlinked process which is begun by the generation of the level-0 data in the case of the simulation study, and which includes a complete homogenisation of the procedures relative to each of its phases in order to guarantee uniformity and therefore comparability of the outputs returned at every step. At the moment, of course the scheme is in an experimental phase and will definitely have to be improved both in terms of the choice of classifiers to be inserted in the set and of the setting of its parameters in order to optimise performances.
Since we are interested in investigating empirically the performances of the single classifiers and of Stacking for datasets with different characteristics and above all if it is convenient to use Stacking in terms of improving performances instead of a single classifier, bearing in mind the necessary increase in computation, some applications of the proposed scheme were carried out both on simulated and real data.
An analysis of the results obtained by applying the Stacking scheme (in the *traditional* approach) to the set of the datasets generated by means of the experimental design does not lead us to believe that the prediction error of the Stacking scheme is to be considered lower than any other classifier

or that, therefore, the Stacking scheme is preferable in terms of performances to the use of the best single classifier.
It always achieves good performances and is to be considered among the best, but it does not seem to be preferable for this type of application, such as Linear Discriminant Analysis and Logistic Regression, while the behaviour of scaled SVM is particularly interesting, as it is has the lowest error among the classifiers, also in comparison with Stacking, LDA and Logistic Regression.
In a set of classifiers in which there was no SVM it would be the best for each of the analysed datasets. In the case of contaminated data Stacking improves its performances noticeably compared to what we have observed for non-contaminated data, in some cases also in comparison with scaled SMV, and generally appears to be very competitive, above all when the contaminations are more substantial and especially in the presence of strong contaminations also for both classes. On the contrary, the effect of the contamination is quite substantial for Linear Discriminant Analysis and Logistic Regression. Generally speaking, based on the outcome of the experiments carried out, it seems to be that the contamination contained by a single class causes a deterioration in the performances of LDA and Logistic Regression, while with a higher level of contamination this only happens with a high degree of separation between the groups. Where there is a low degree of separation, or both classes are contaminated, the worst performance will be achieved for all four of the parametric classifiers.
A further element of interest in our research was whether the number and typologies of the algorithms chosen were important for the composition of the set of base classifiers, since the literature is quite unequivocal in maintaining that any classifier can be used in a Stacking scheme. Different Stacking schemes have been created for different input datasets. The main results obtained from the application of the different Stacking schemes show that the choice from the start of high performance base-level classifiers could increase the predictive capacities of Stacking, and however, a comparison shows that the Stacking scheme with three base classifiers, is more successful (although the differences are slight). There would seem to be a confirmation of the circumstance that a greater complexity of the meta model does not improve results also for the applications to real data : in fact the best Stacking performances were achieved (at least in the examples analysed) with schemes with a lower number of base classifiers STA3 and STA6.
As far as the application to real data is concerned, the analysis was carried out on two datasets, the first was particularly suitable for classification, while the second is a very complex dataset containing the results of the inspection surveys carried out by INPS (National Social Security Institute) on Italian companies in order to see if there was any off-the-book employment present, and it achieves error values that are generally quite high for all classifiers, as was predictable given the level of overlap between the groups. It is easy to guess from the structure of the data that there are anomalous observations which are in the group represented by the companies that declare an absence of any off-the-book employment. Any anomalous values could give rise to signals of the potential presence of off-the-book employment in those companies. In both cases, however, Stacking seems competitive when compared to the use of ensemble methods (Bagging and Adaboost), but not preferable in terms of performances to the use of the best single classifier apart the STA 6 scheme.

Moving on to the most innovative element of the thesis, that is the extension of the Forward Search approach to the Stacking scheme it seems quite clear that in the passage from the traditional approach to Forward Search, we acquire a much wider knowledge of the structure of the data, the trajectories of the observations, the importance exerted by each of these ones, whether outlier or not, on the model or on the aspects that are of interest for evaluating the performances of the base classifiers and of the Stacking scheme. With the traditional modality we could only see a static result, which is equivalent to the last step of the search but no information about the effect exerted by the single units, since they were treated simultaneously in the analysis, even if there was a logic of sample rotation such as cross-validation.
By starting from an observation of the plots of Mahalanobis distances, which show the entrance of the outliers (or influential observations) graphically, we can evaluate the alteration of the performances of the classifiers, or rather the distortion of the classification rule by verifying at the entrance point the effect caused by atypical observations in the plot of cross-validation errors for units inside the subset and in the plots of the posterior and cross-validation probabilities that show the probabilities for each unit of belonging to the relative group at every step of the search.

From the analysis of the results of some applications of the Stacking scheme proposed in a Forward Search context it seems quite clear from the plots of the posteriors and of the cross validation posterior probabilities referred to the simulated and real data, that the classifiers display different behaviour also in the case where there are no outliers.
In the case of simulated contaminated data, the plots of the minimum and maximum Mahalanobis distances clearly reveals the entrance of anomalous observations into the subset, and, in the cross validation error graph, at the point where we have established that anomalous observations make their entrance, there is a general increase in the curves relative to classifiers but there is clearly an increase for those relative to Linear Discriminant Analysis and Logistic Regression (while the contamination of two classes has decreased the performances of the four parametric classifiers). Unlike what we have seen in the traditional analysis, we received very important additional information: in fact, we are able to see in exactly which moment the decision rule began to decline, that is when the anomalous values make their entrance and therefore to monitor the decision rule, which is not possible in the traditional approach. This has allowed us also to see that Stacking was not only the best classifier but it also remained the best for all the steps of the search, thus highlighting the stability of its classification rule.
In the first application to real data we were interested to explore data and investigate the behaviour of the single classifiers and of Stacking and above all to see there is an accordance between the fit of the model and prediction. Therefore, in terms of full allocation, in the balanced search there are no significant differences between posterior and cross validation with the exception of Adaboost and Bagged Classification Tree for which the model is in the prediction phase, able to fully allocate (probability ≥ 0.8) a number of units higher than that achieved by means of posterior probabilities, while in the unbalanced search, this is only the case for Adaboost. There is, however, a difference in terms of unit allocation (and therefore of the performance of the classifier) and in the posterior probabilities with the exception of Adaboost all the units are allocated to the relative group, even if belatedly, at the last step (as for unit 9 in the first group) and even if it is not done perfectly. The situation is not the same for cross validation probabilities. Only LDA, QDA and Stacking will

allocate all the units to the respective groups even though not perfectly (LDA) or in the last steps both in balanced and in unbalanced search.
The last two applications are referred to datasets for which a transformation of the data is necessary given the structure of the data and for which we used the automatic outlier detection procedure proposed in Riani, Atkinson and Cerioli (2009), even if this was done separately for each population.

Since the main results obtained from the application of the proposed automatic outlier detection procedure, even if separately for each population, appear to be very interesting, as far as more in-depth studies and future developments in research, the extension in a multivariate context for problems of supervised classification.

With a view to a further improvement of the entire proposed process, which is at the experimental stage, the research activity will be directed towards optimising performances and guaranteeing the reliability of the predictions for single classifiers by modifying the setting of the parameters used in this phase, and more generally by including:

- Extension of the experimental design both to verify further the results achieved and to insert other elements into the design (different processes of data generation, increasing the number of classes, different *prior* values)
- Introduction of more adequate measurements (at some point combining the use of more than one index) which are able to capture accuracy in the best possible way in terms of estimating the prediction error returned by the classifiers
- Possible introduction of a weighting system into the method of meta-classification should we intend to combine several classifiers with very different performances in terms of accuracy.
- Extension of the methods proposed by the literature (Varma et al., 2006 ; Tibshirani et al., 2008) for the estimation and reduction of potential *bias* in cross validation error for the problem that is the object of the thesis.

APPENDIX to Chapter 3

This Appendix gives the traditional formalization of the four parametric algorithms used to induce the respective base classifiers. We recall, therefore, some notations and concepts commonly used for dealing with classification problems

- **Linear Discriminant Analysis**

In a $k-class$ classification problem we need to know the class posterior probabilities $\Pr(G \mid X)$ for optimal classification, where X is a casual p-dimensional variable, and G is a casual categorical variable that represents the to which an individual belongs.
The overall population is made up of K classes and we suppose $f_k(x)$ is the class-conditional density of X in class $G = k$, and let π_k be the prior probability of class k, with $\sum_{k=1}^{k} \pi_k = 1$.
Therefore, the density for the overall probability is:

$$f_k(x) = \sum_{k=1}^{K} f_k(x)\pi_k .$$

By application of the Bayes theorem:

$$\Pr(G = k / X = x) = \frac{f_k(x)\pi_k}{\sum_{\ell=1}^{k} f_\ell(x)\pi_\ell} \qquad (1.1)$$

We will have to know or evaluate from the data π_k and $f_k(x)$. In a parametric context we can suppose the hypothesis that we model each class density as multivariate Gaussian $N_p(\mu_k, \Sigma_k)$, so that the result is

$$f_k(x) = \frac{1}{(2\pi)^{p/2}\left|\Sigma_k\right|^{1/2}} e^{-\frac{1}{2}(x-\mu_k)^T \Sigma_k^{-1}(x-\mu_k)} \qquad \text{for } k = 1,\ldots,K \qquad (1.2)$$

Linear discriminant analysis (LDA) arises in the special case when we assume that the classes have a common covariance matrix $\sum_k = \forall k$

The *linear discriminant functions (LDA)*

$$\delta_k(x) = x^T \Sigma^{-1} \mu_k - \frac{1}{2}\mu_k^T \Sigma^{-1} + \log \pi_k \qquad (1.3)$$

are an equivalent description of the decision rule, with $G(x) = \arg\max_k \delta_k(x)$. Generally, we do not know the parameters of the Gaussian distributions. We will need to estimate them from the data as follows:

- $\hat{\pi}_k = n_k / n$
- $\hat{\mu}_k = \Sigma_{g_i=k} x_i / n_k$
- $\hat{\Sigma} = \sum_{k=1}^{K} \sum_{g_i=k} \left(x_i - \hat{\mu}_k\right)\left(x_i - \hat{\mu}_k\right)^T / \left(n - K\right)$

For a two-class problem there is a correspondence between Linear Discriminant Analysis and classification by linear least squares. The decision rule is assigned to class 2 if

$$x^T \hat{\Sigma}^{-1}\left(\hat{\mu}_2 - \hat{\mu}_1\right) > \frac{1}{2}\hat{\mu}_2^T \hat{\Sigma}^{-1}\hat{\mu}_2 - \frac{1}{2}\hat{\mu}_1^T \hat{\Sigma}^{-1}\hat{\mu}_1 + \log\left(n_1 / n\right) - \log\left(n_2 / n\right)$$

and class 1 otherwise. Suppose we code the targets in the two classes as *+1* an *-1* respectively. With more than two classes, LDA is not the same as linear regression of the class indicator matrix, and it avoids the masking problems associated with that approach (Hastie et al., 1994).

- **Quadratic Discriminant Analysis**

If, on the other hand, the Σ_k are not assumed to be equal, from the expression of density (1.2) we then get *quadratic discriminant functions* (QDA),

$$\delta_k(x) = \frac{1}{2}\log|\Sigma_k| - \frac{1}{2}\left(x - \mu_k\right)^T \Sigma_k^{-1}(x - \mu_k) + \log \pi_k \qquad (1.4)$$

The estimates for μ_k for QDA are similar to those for LDA, while separate covariance matrices must be estimated for each class:

$$\hat{\Sigma}_k = \Sigma_{g_i=k}\left(x_i - \hat{\mu}_k\right)\left(x_i - \hat{\mu}_k\right)^T / \left(n_k - 1\right)$$

Unlike the Linear Discriminant Analysis, QDA is closely linked to the Gaussian distributive hypothesis.

- **Naïve Bayes**

Naïve Bayes models are a variant of the previous case, and assume that each of the class densities are products of marginal densities; that is, they assume that given a class $G = j$, the features X_k are independent:

$$f_j(X) = \prod_{k=1}^{p} f_{jk}(X_k)$$

This assumption is often not true, but it simplifies the estimation. The individual class-conditional marginal densities f_{jk} can each be estimated separately using one-dimensional kernel density estimates. This is in fact a generalization of the original *Naïve Bayes* procedures,which used univariate Gaussians to represent these marginals. Naïve Bayes uses a Laplacian estimate for estimating the conditional probabilities for each nominal attribute to compute f_{jk} . For each continuous-valued attribute, a normal distribution is assumed in which case the conditional probabilities can be conveniently represented entirely in terms of the mean and variance of the observed values for each class.

- **Logistic Regression**

The specific form of the Logistic Regression model for the posterior probabilities $P_k(x)$ via linear functions in *x*, while at the same time ensuring that they sum to one and remain in [0, 1] if there are two classes is:

$$P(G=1 \mid X=x) = \frac{e^{(\beta_0+\beta^T x)}}{1+e^{(\beta_0+\beta^T x)}}$$

$$P(G=2 \mid X=x) = \frac{e^{(\beta_0+\beta^T x)}}{1+e^{(\beta_0+\beta^T x)}}$$

The *logit* transformation in terms of $P_k(x)$ is defined : $g(x) = \log\left[\frac{p}{(1-p)}\right]$ and the model has the form:

$$\log \frac{P(G=1 \mid X=x)}{P(G=2 \mid X=x)} = \beta_0 + \beta^T x$$

The importance of this transformation is that $g(x)$ has many of the properties of a linear regression model. The logit $g(x)$ is linear in its parameters, may be continuous, and may range $-\infty$ to $+\infty$, depending on the range of *x*.

- **Bagging** (Bootstrap aggregation)

Given a learning set $L = \{(y_n, x_n), n = 1, ..., N\}$, where the *y* are class labels, we assume to build a model on it obtaining the prediction $\hat{f}(x)$ at input *x*.
Bootstrap aggregation (or *Bagging)* averages this prediction over a collection of bootstrap samples $\{L^{(B)}\}$ from L.

For each bootstrap sample $L^{(B)}$ we fit our model, that returns prediction $\hat{f}^{(B)}(x)$.
The bagging estimate is defined by:

$$\hat{f}_{bag}(x) = \frac{1}{B}\sum_{b=1}^{B} \hat{f}^{(B)}(x)$$

If *y* is a class label, let the $\hat{f}^{(B)}(x)$ vote to form $\hat{f}_{bag}(x)$.
The $\{L^{(B)}\}$ form replicate data sets, each consisting of *N* cases, drawn at random, *but with replacement,* from *L*. Each (y_n, x_n) may appear repeated times or not at all in any particular $L^{(B)}$.
The $\{L^{(B)}\}$ are replicate data sets drawn from the bootstrap distribution approximating the distribution underlying L.

A critical factor in whether bagging will improve accuracy is the stability of the procedure for constructing $\hat{f}$. If changes in *L*, i.e. a replicate *L*, produces small changes in $\hat{f}$, then $\hat{f}_{bag}$ *B* will be close to $\hat{f}$. Improvement will occur for unstable procedures where a small change in L can result in large changes in $\hat{f}$. (Breiman 1996).

- **AdaBoost**

Boosting is a general method for improving the performance of any learning algorithm. Schapire introduced the first boosting algorithm in 1990. In 1995, Freund and Schapire introduced the AdaBoost algorithm. In this thesis, we refer to AdaBoostM1 (Freund and Schapire,1996).

The algorithm assumes a training set consisting of *m* instances $S = \{(x_1, y_1), ..., (x_m, y_m)\}$ where x_i is a vector of attribute values and $y_i \in Y$ is the class label associated with x_i. The boosting algorithm call another unspecified learning algorithm (called WeakLearn) repeatedly in a series

of rounds. The purpose of the boosting is to apply the weak learner to repeatedly modified version of the data, producing a sequence of weak classifiers h_t and the predictions from all of them combined through a weighted majority vote to give the final prediction that minimizes the error. On round *t*, therefore, the booster provides WeakLearn with a distribution D_t over the training set S and in response it computes a classifier which should correctly classify a fraction of the training set that has large probability with respect to . The process is carried out for $t = 1, 2, ..., T$ and in T the booster combines all weak classifiers into a final classifier h_{fin}.

Algorithm AdaBoost.M1

Input: $S = \{(x_1, y_1), ..., (x_m, y_m)\}$

WeakLearn

T (number of iterations)

Inizialize $D_1(i) = 1/m; i = 1, ..., m$

For t =1 to T

1. Fit a classifier h_t using WeakLearn and distribution D_t
2. Compute error of h_t: $\varepsilon_t = \sum_{i:h_t(x_i) \neq y_i} D_t(i)$

   If $\varepsilon > 0.5$ then $T \leftarrow t-1$ exit Loop
3. $\beta_t = \varepsilon_t / (1-\varepsilon_t)$
4. Update distribution D_t

   $$D_{t+1}(i) = \frac{D_t(i)}{Z_t} \times \begin{cases} \beta_t & if\ h_t(x_i) = y_i \\ 1 & otherwise \end{cases}$$

   Z_t is normalization constant in order to be D_{t+1} a distribution

Output the final classifier

$$h_{fin}(x) = \arg\max_{y \in Y} \sum_{i:h_t(x_i)=y_i} \log \frac{1}{\beta_t}$$

APPENDIX to Chapter 4

	N=120				N=200			
	δ=0.5		δ=3		δ=0.5		δ=3	
	n.variables				n.variables			
Classifier	3	10	3	10	3	10	3	10
BAG	0,3879	0,2821	0,0188	0,0018	0,3911	0,2770	0,0158	0,0011
ADA	0,3836	0,2810	0,0223	0,0218	0,3730	0,2717	0,0181	0,0152
SVMb	0,3378	0,2256	0,0044	0	0,3354	0,2236	0,0044	0
STA	**0,3547**	**0,2323**	**0,0048**	**0**	**0,3498**	**0,2268**	**0,0058**	**0**

Table A1. Cross validation error for three base classifier and STA3 scheme for input datasets with varying degrees of complexity and with different degrees of separation between the two groups. N=120. N=200

	N=120				N=200			
	δ=0.5		δ=3		δ=0.5		δ=3	
	n.variables				n.variables			
Classifier	3	10	3	10	3	10	3	10
TRE	0,4202	0,3617	0,0372	0,0416	0,4174	0,3554	0,0330	0,0326
TREnp	0,4098	0,3624	0,0391	0,0401	0,4141	0,3553	0,0313	0,0303
BAG	0,3893	0,2832	0,0193	0,0017	0,3898	0,2794	0,0144	0,0007
ADA	0,3798	0,2833	0,0208	0,0215	0,3741	0,2736	0,0162	0,0155
STA	**0,3903**	**0,2798**	**0,0239**	**0,0059**	**0,3801**	**0,2732**	**0,0168**	**0,0025**

Table A2. Cross validation error for four base classifier and STA4 scheme for input datasets with varying degrees of complexity and with different degrees of separation between the two groups. N=120. N=200

	N=120				N=200			
	δ=0.5		δ=3		δ=0.5		δ=3	
	n.variables				n.variables			
Classifier	3	10	3	10	3	10	3	10
LDA	0,3413	0,2416	0,0053	0,0000	0,3417	0,2326	0,0009	0
QDA	0,3527	0,2974	0,0053	0,0000	0,3523	0,2696	0,0016	0
BAG	0,3912	0,2857	0,0190	0,0014	0,3893	0,2764	0,0076	0,0008
ADA	0,3792	0,2808	0,0208	0,0209	0,3743	0,2761	0,0324	0,0156
NBA	0,3473	0,2419	0,0051	0	0,3466	0,2339	0,0008	0
GLM	0,3420	0,2417	0,0093	0	0,3411	0,2330	0,0023	0
STA	**0,3472**	**0,2399**	**0,0066**	**0**	**0,3450**	**0,2324**	**0,0016**	**0**

Table A3. Cross validation error for six base classifier and STA6 scheme for input datasets with varying degrees of complexity and with different degrees of separation between the two groups. N=120. N=200

Appendix to Chapter 6

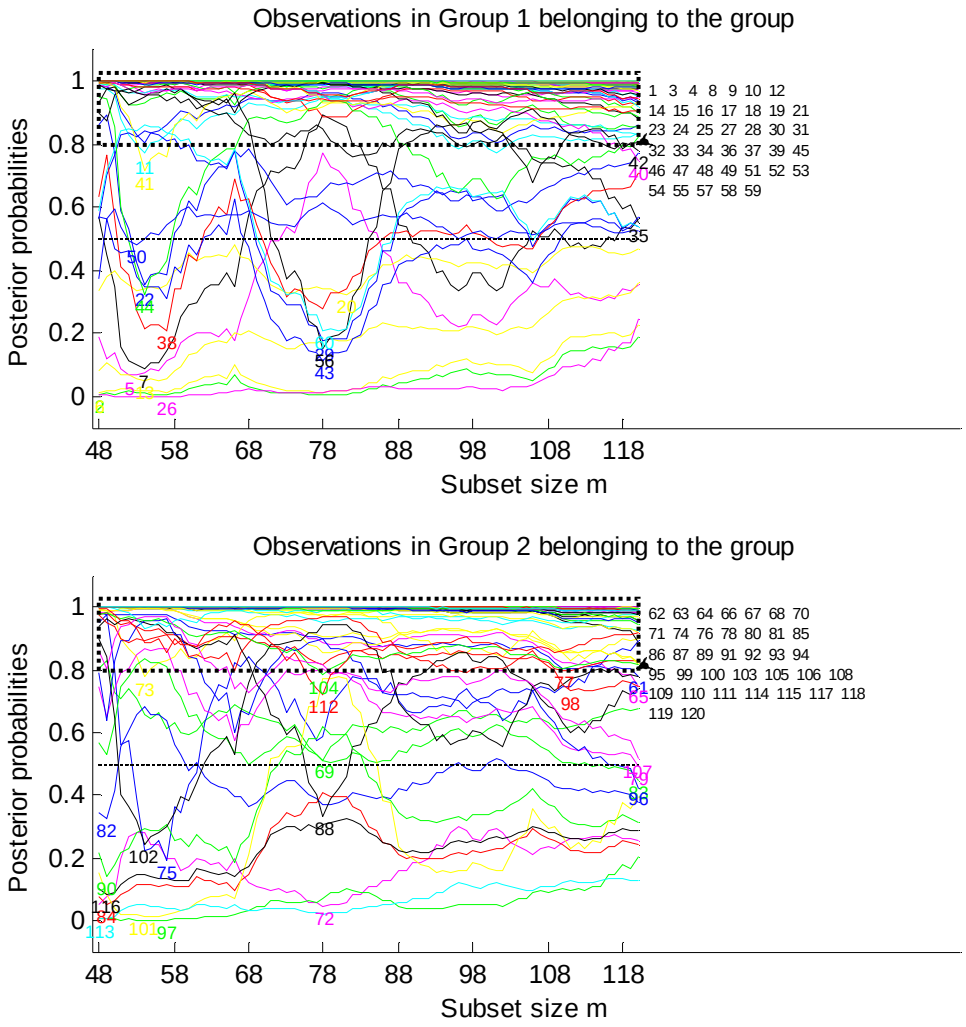

Figure B1. Simulated data. Linear Discriminant Analysis: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.

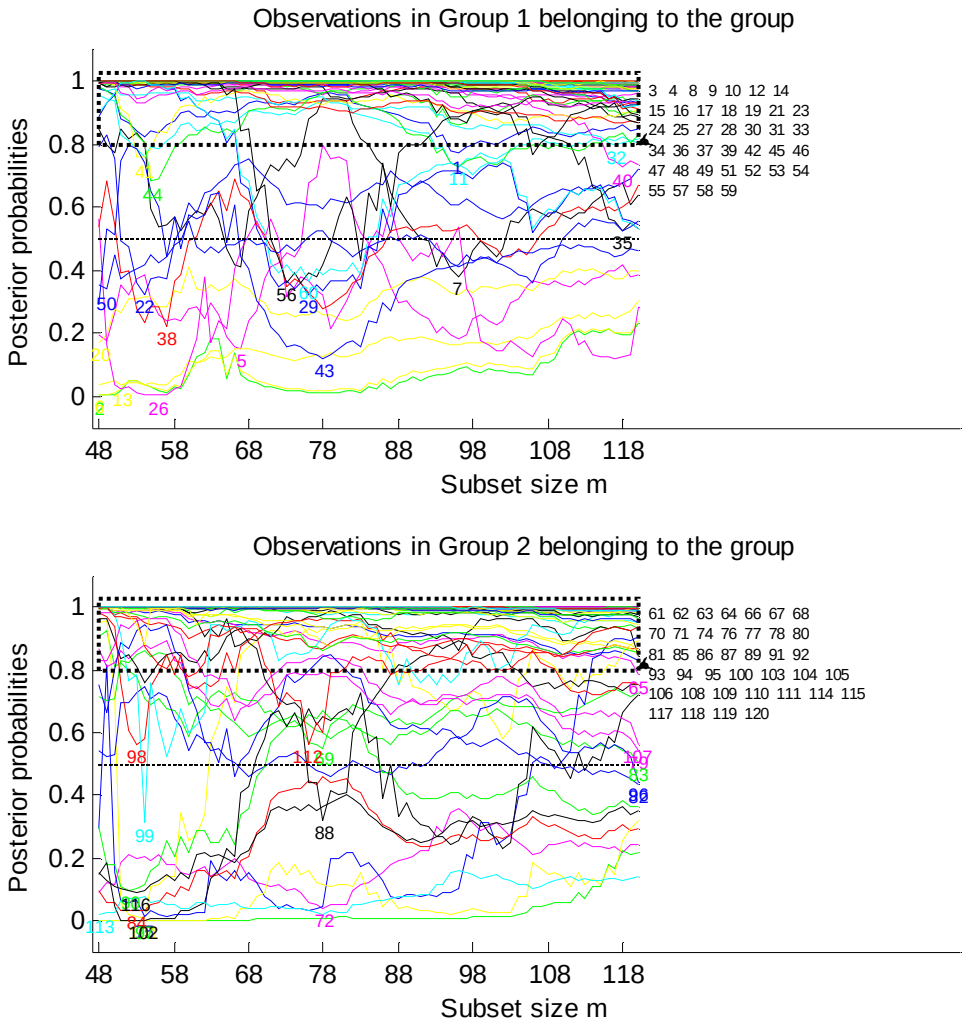

Figure B2. Simulated data. Quadratic Discriminant Analysis: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.

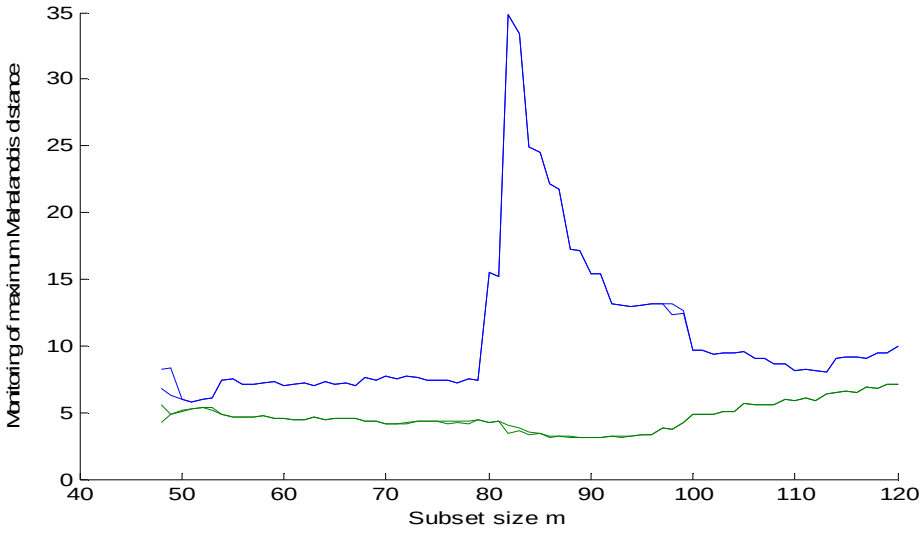

Figure B3. Simulated contaminated data. Plot of maximum Mahalanobis distance for units inside the subset. Balanced search.

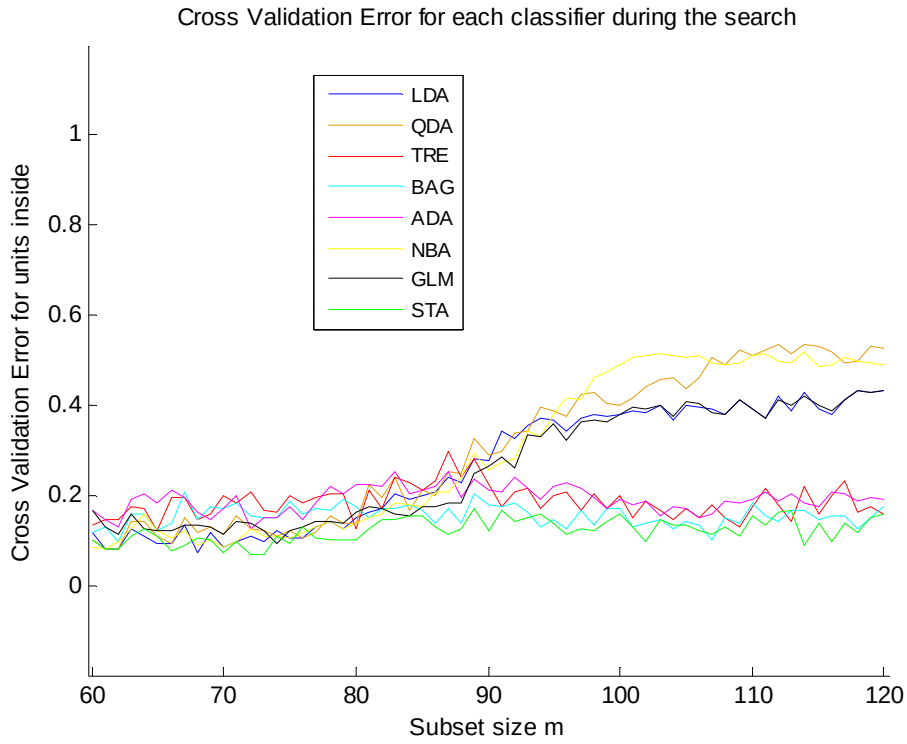

Figure B4. Simulated contaminated data. Plot of cross validation error for units belonging to the subset of the seven base classifiers and Stacking. Balanced search. The contamination is intended for two class.

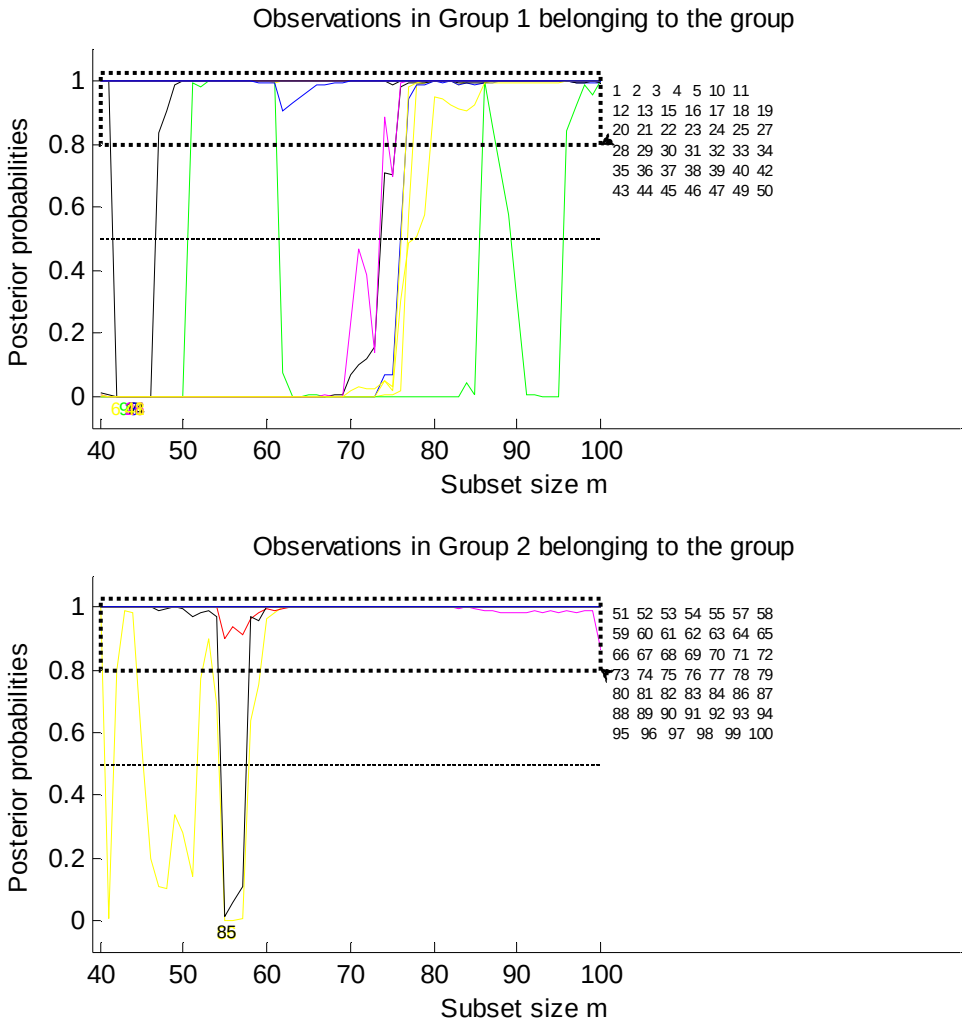

Figure B5. Electrodes data. Quadratic Discriminant Analysis: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

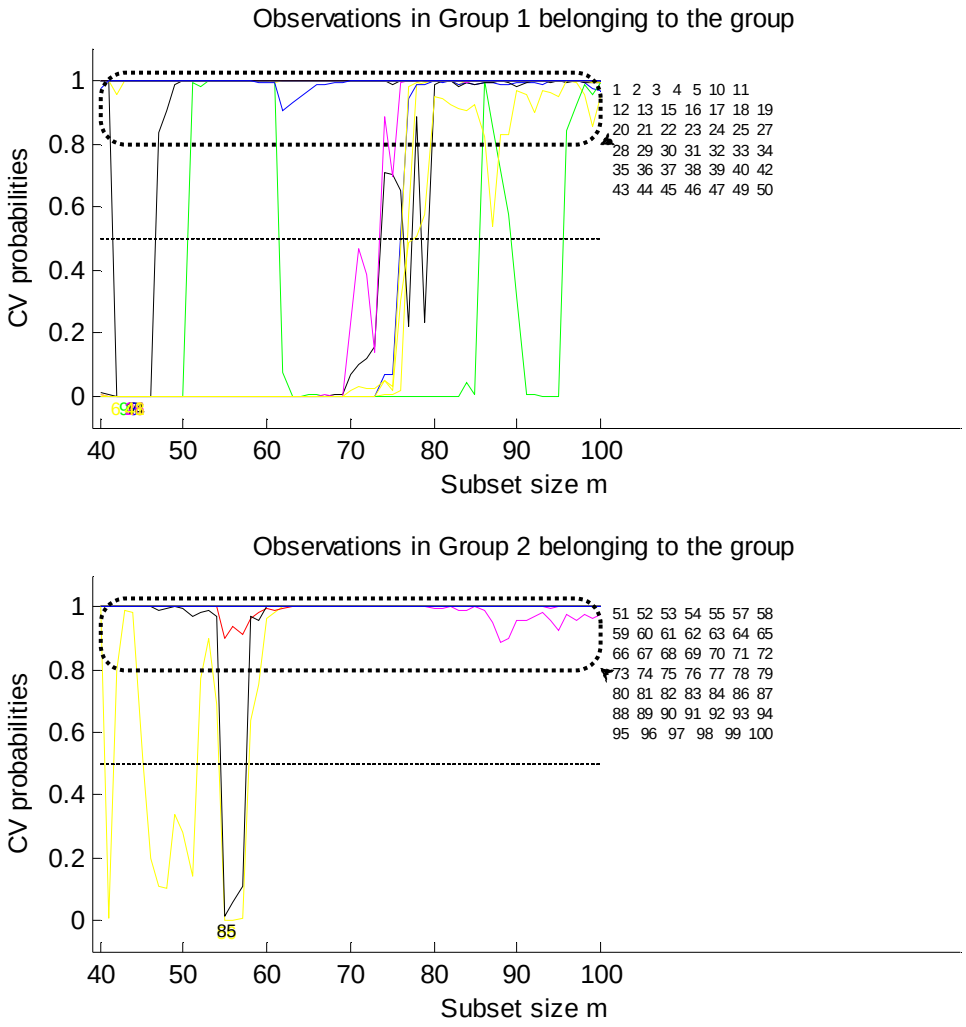

Figure B6. Electrodes data. Quadratic Discriminant Analysis: cross validation error probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

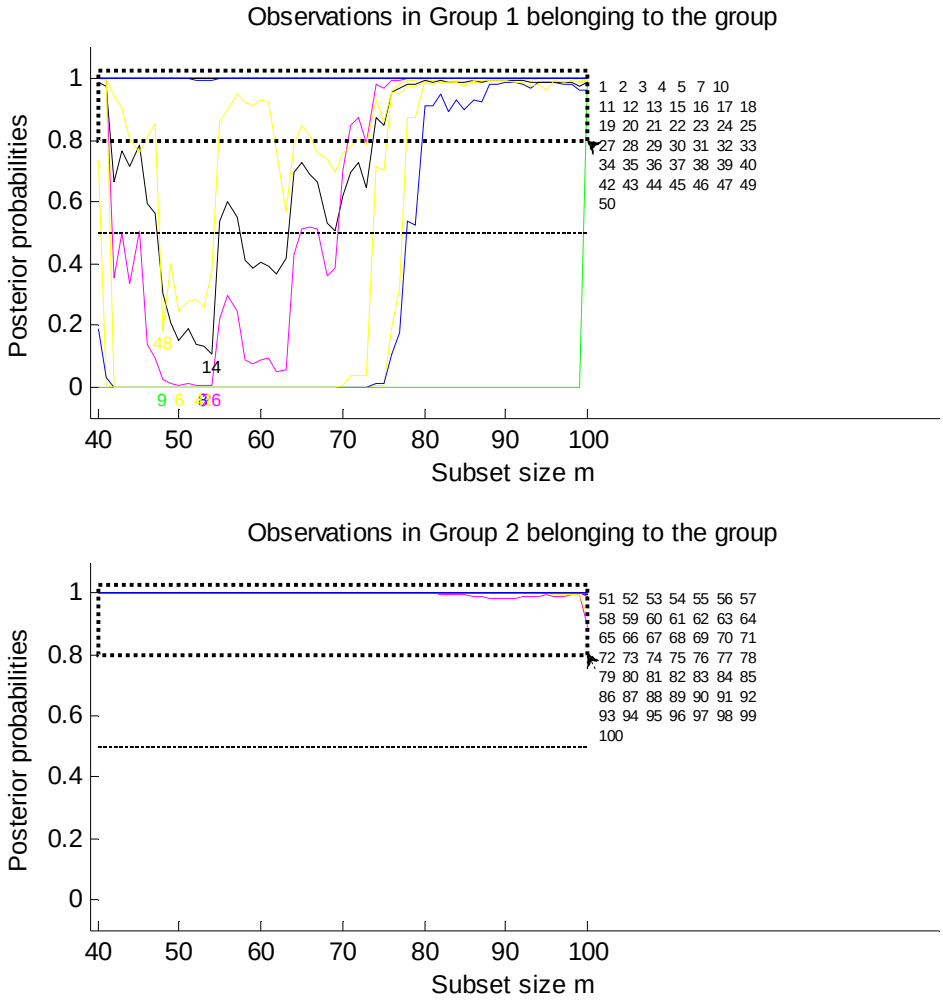

Figure B7. Electrodes data. Naive Bayes: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

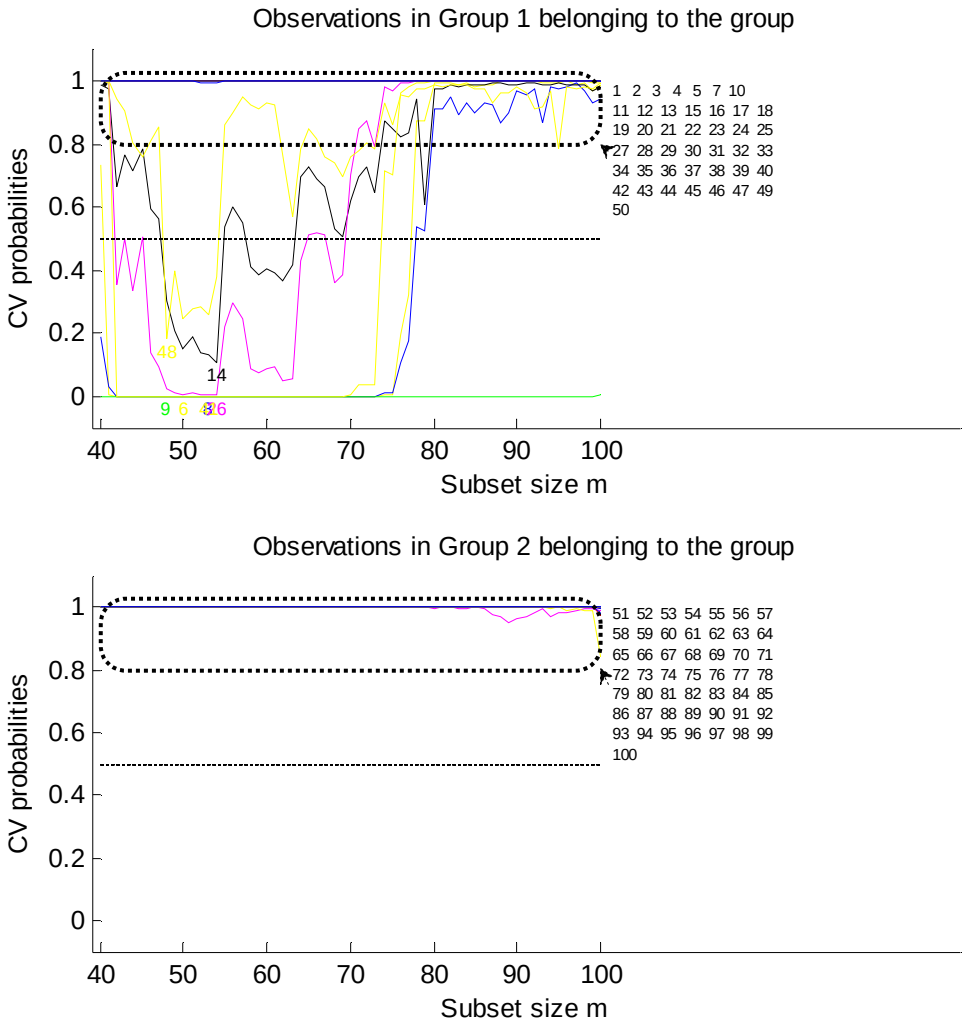

Figure B8. Electrodes data. Naive Bayes: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

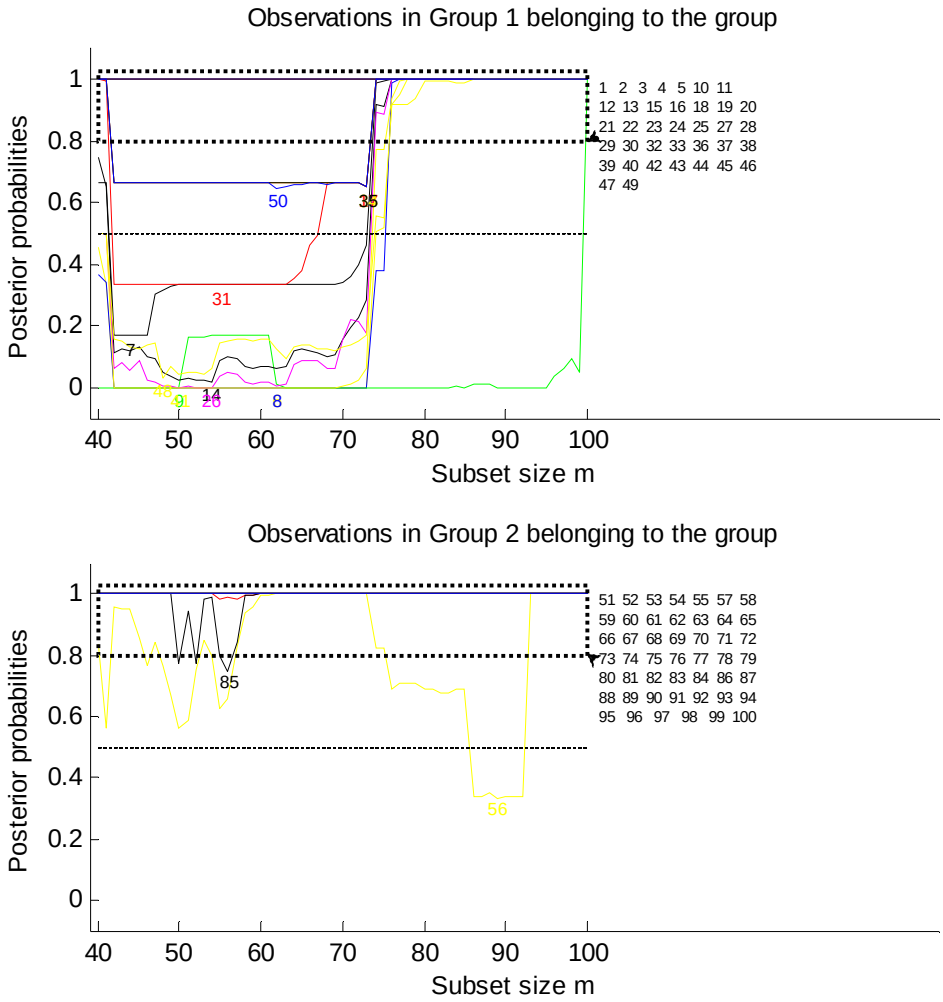

Figure B9. Electrodes data. Stacking scheme: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

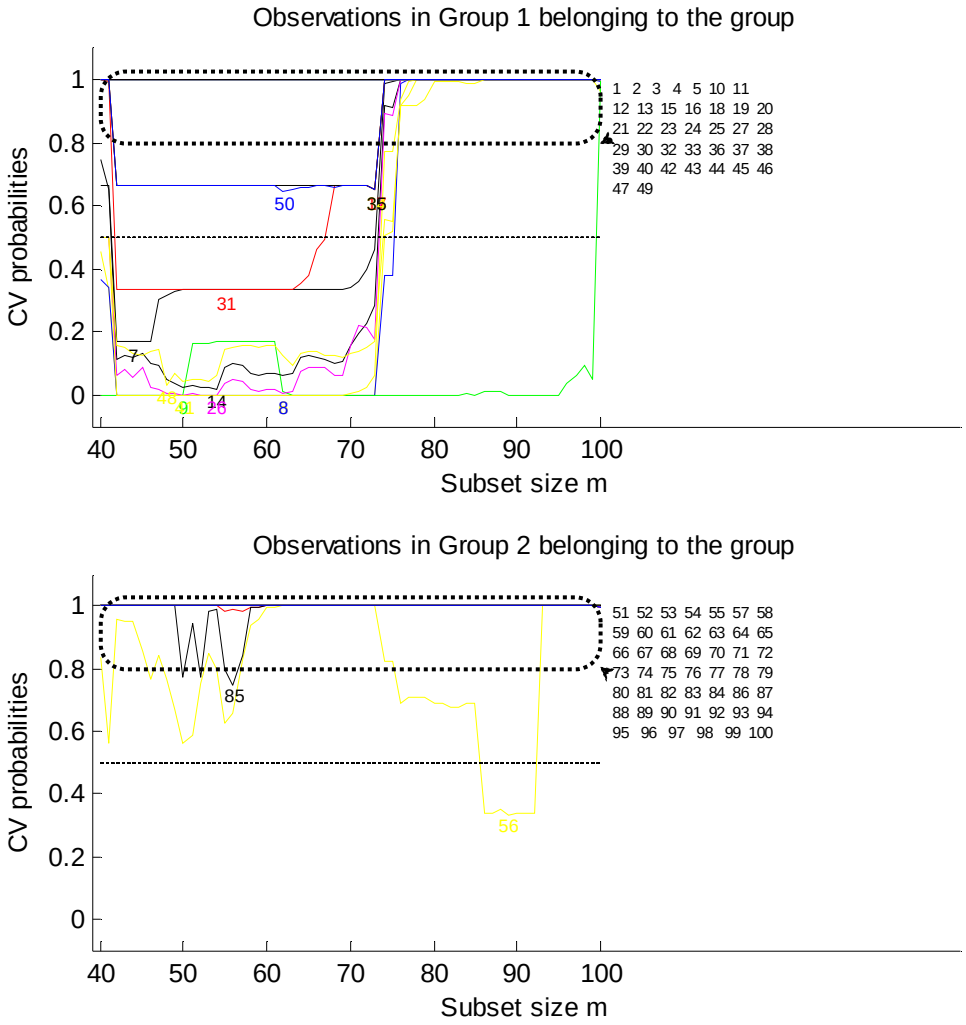

Figure B10. Electrodes data. Stacking scheme: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

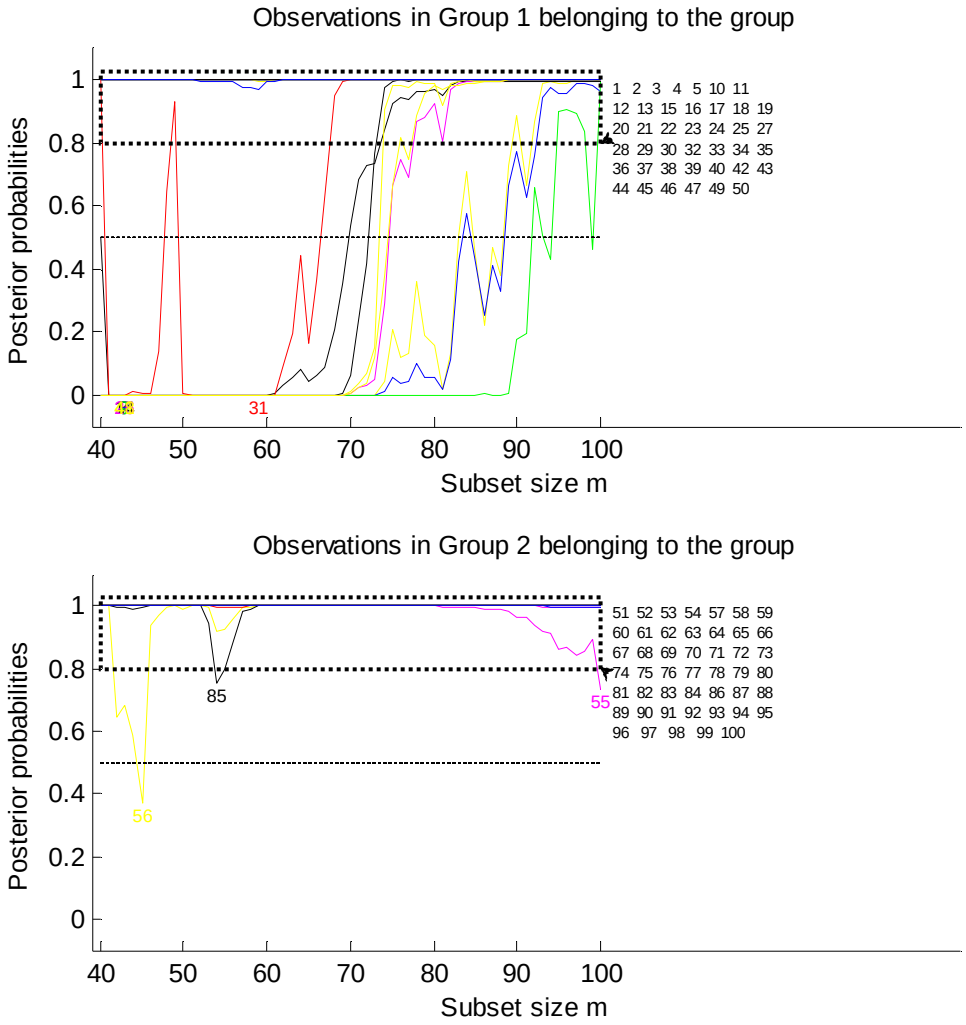

Figure B11. Electrodes data. Linear Discriminant Analysis: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unbalanced search

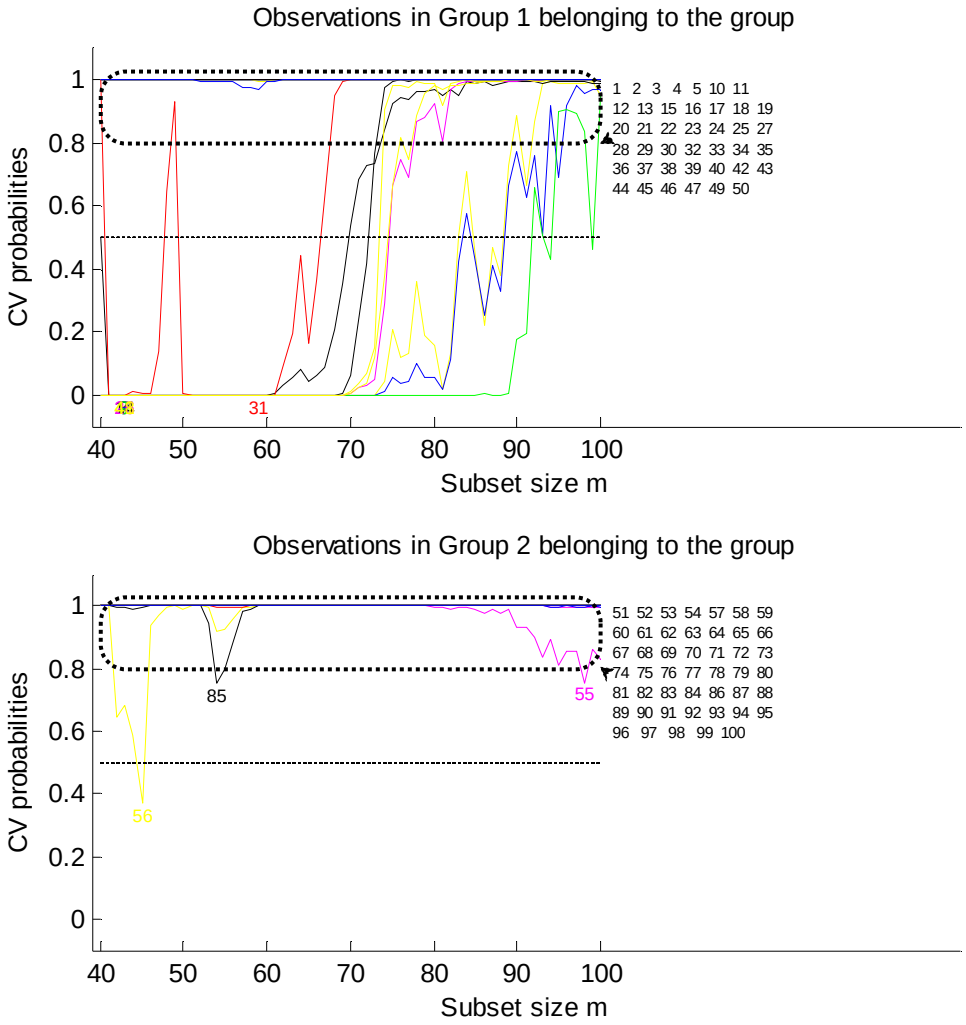

Figure B12. Electrodes data. Linear Discriminant Analysis: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unalanced search

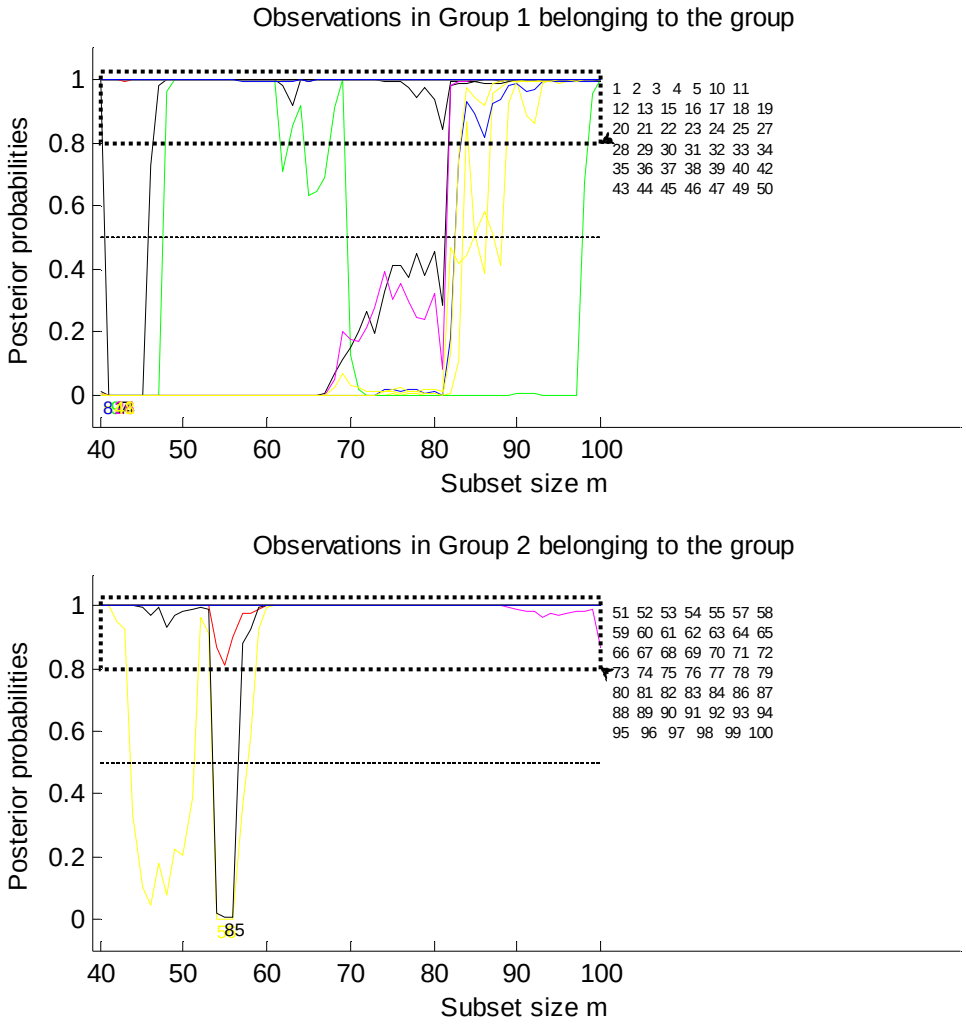

Figure B13. Electrodes data. Quadratic Discriminant Analysis: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unbalanced search

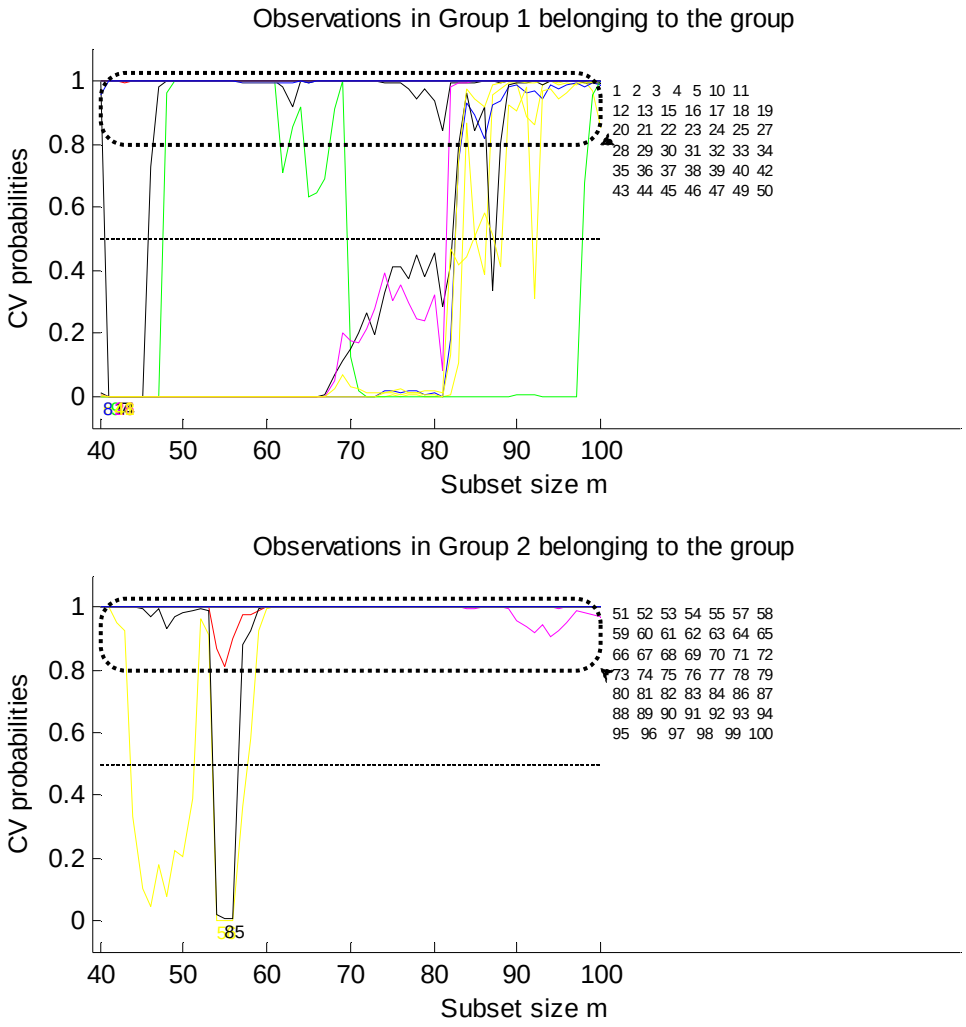

Figure B14. Electrodes data. Quadratic Discriminant Analysis: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unbalanced search

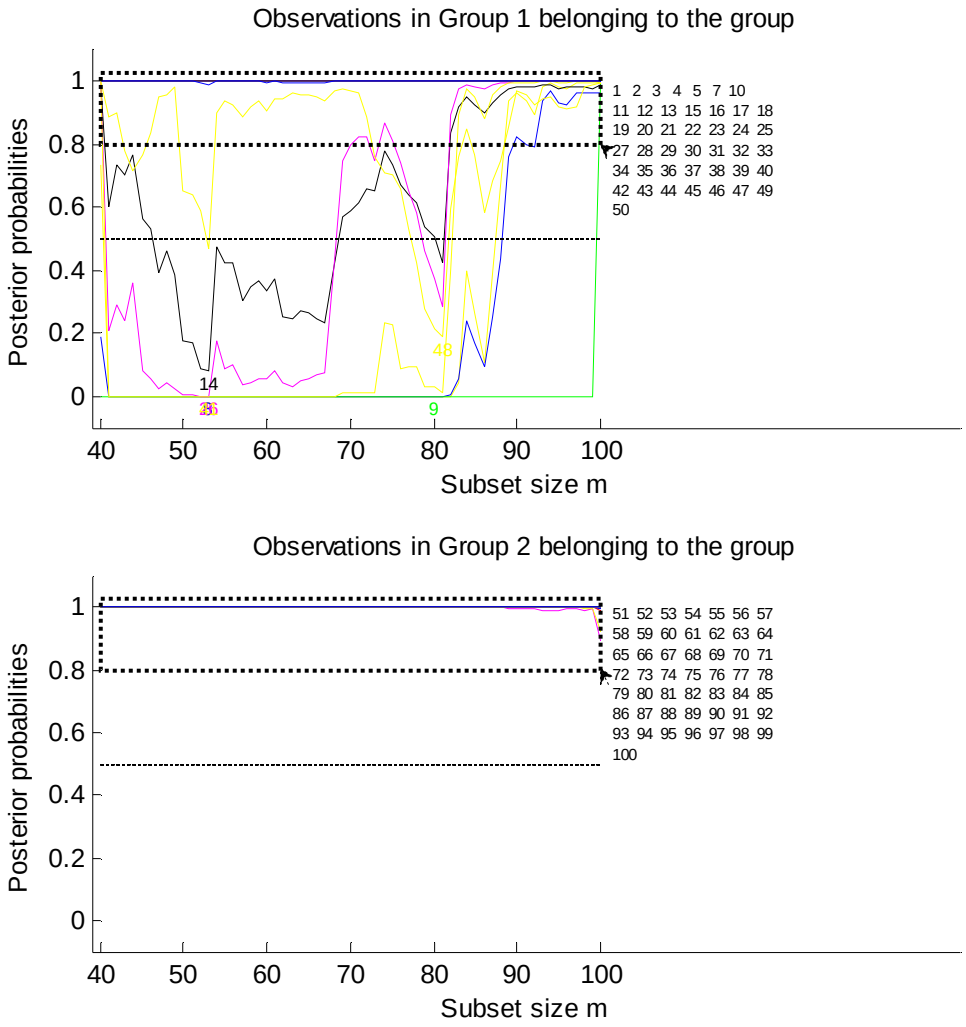

Figure B15. Electrodes data. Naive Bayes: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unbalanced search.

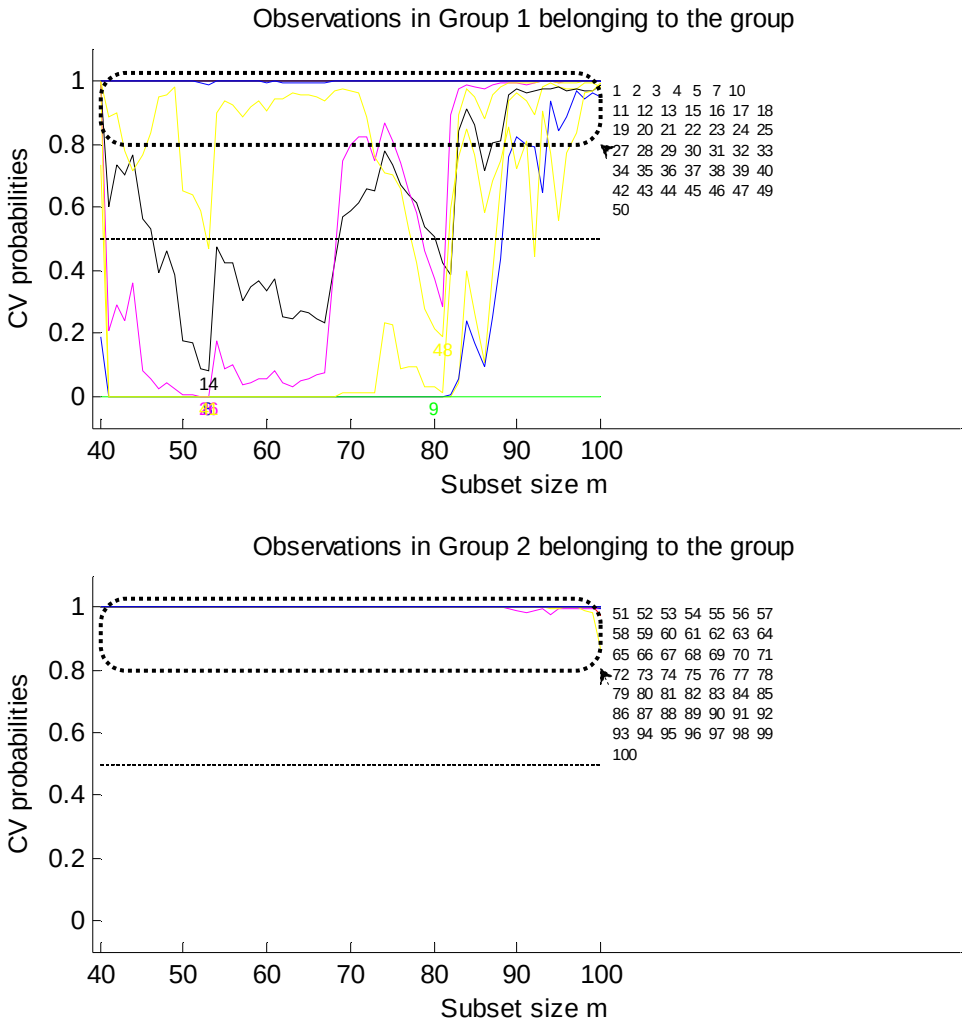

Figure B16. Electrodes data. Naive Bayes: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unbalanced search.

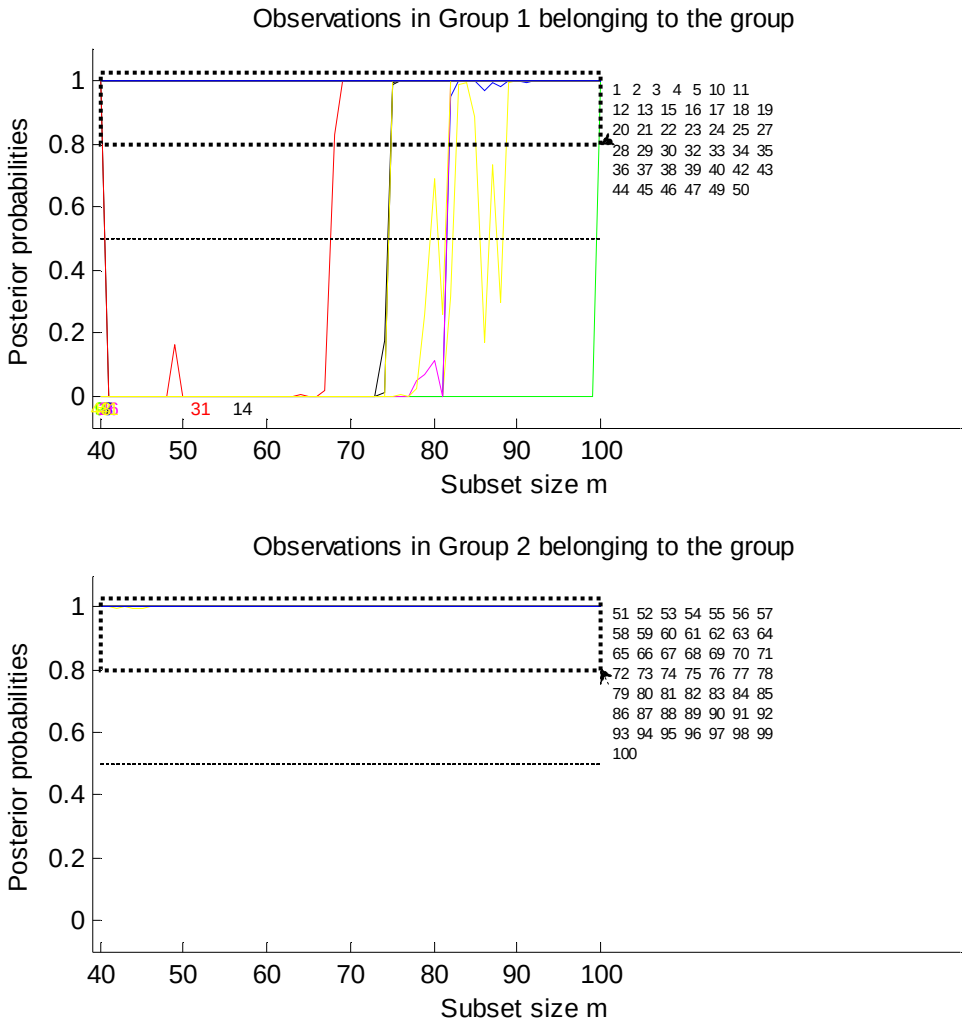

Figure B17. Electrodes data. Logistic Regression: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unbalanced search.

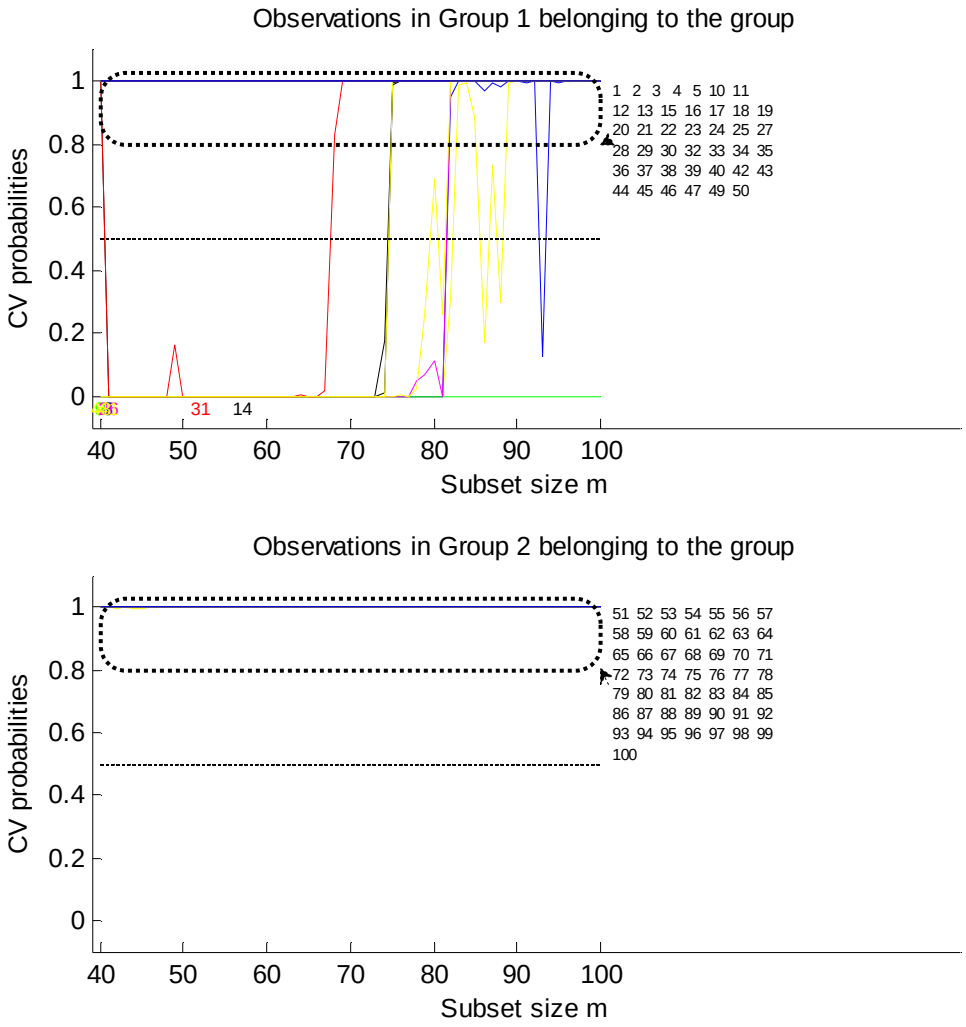

Figure B18. Electrodes data. Logistic Regression: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unbalanced search.

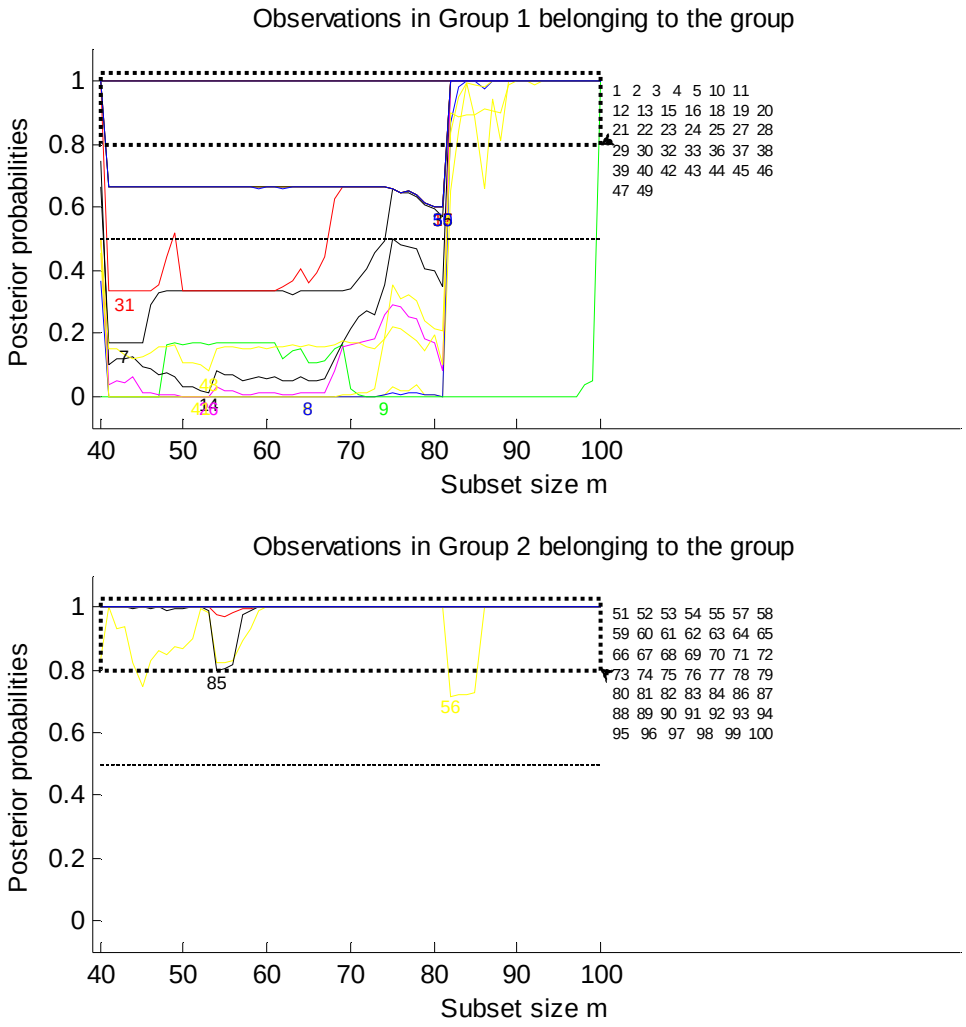

Figure B19. Electrodes data. Stacking scheme: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unbalanced search.

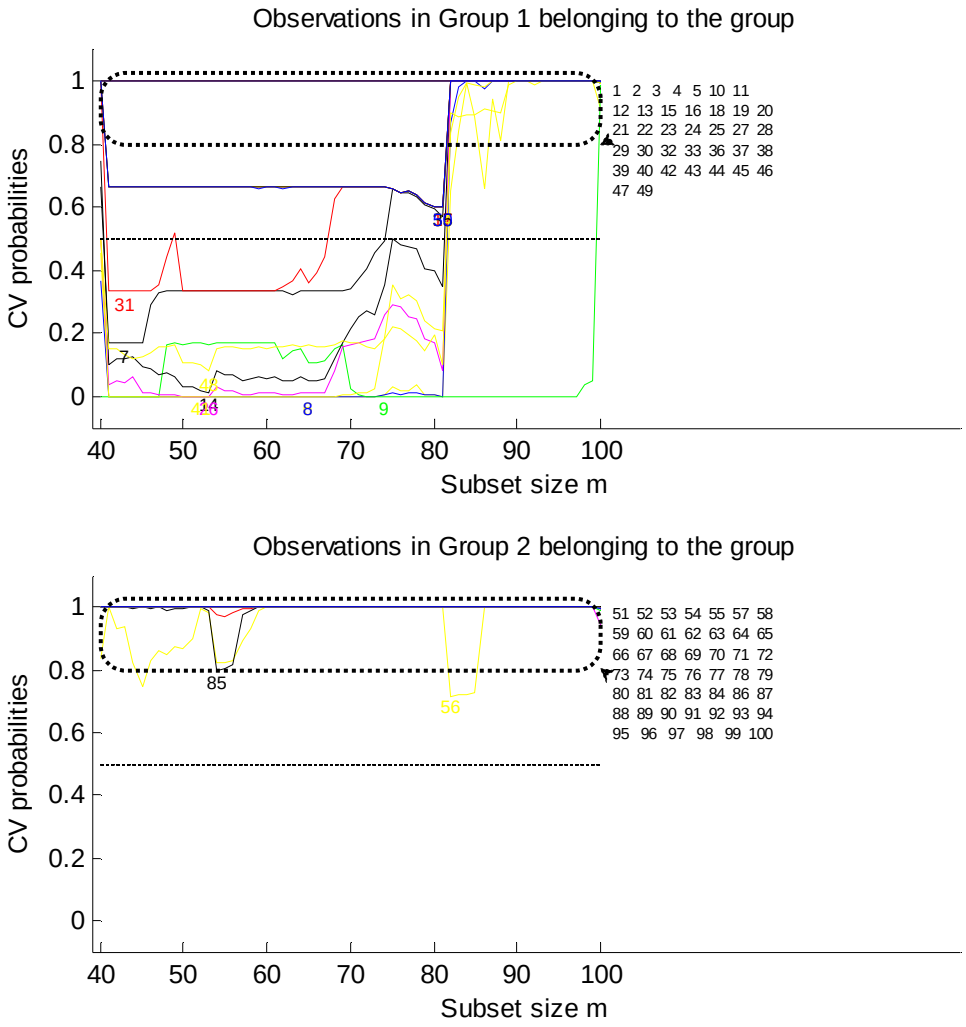

Figure B20. Electrodes data. Stacking scheme: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Unbalanced search.

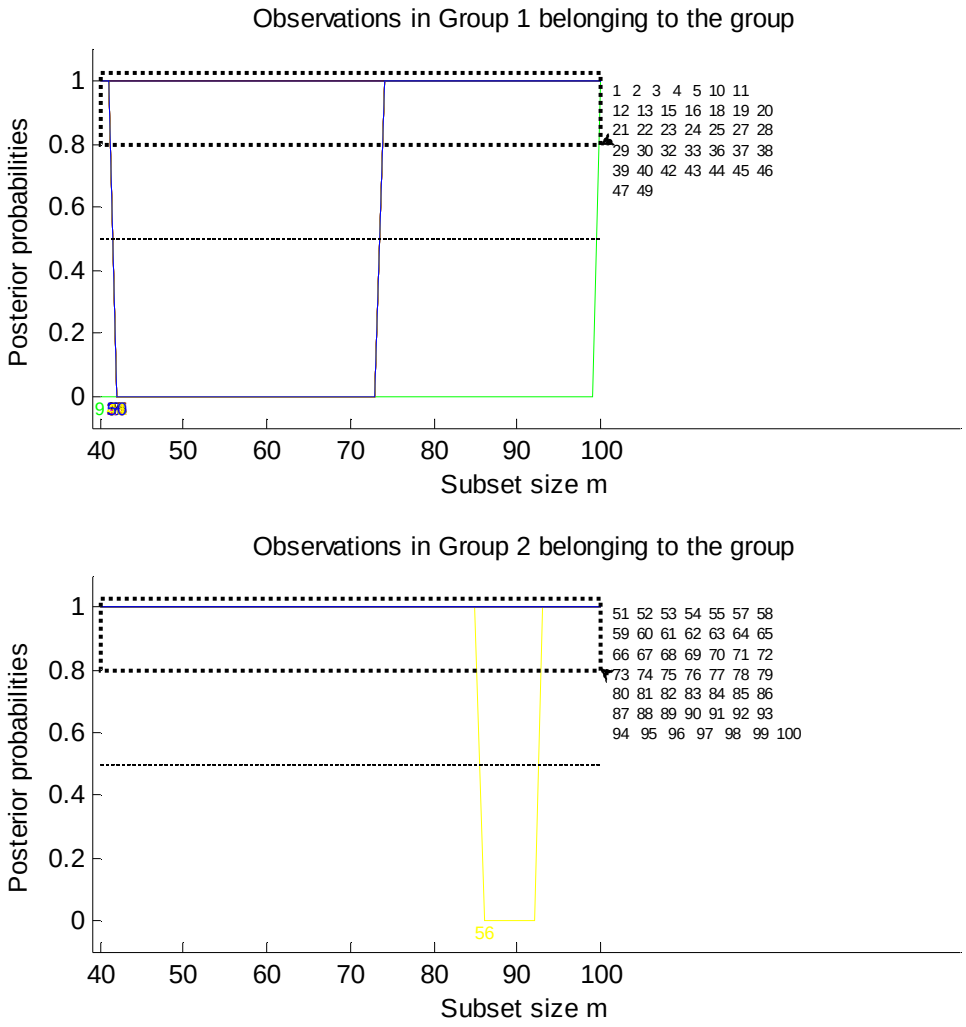

Figure B21. Electrodes data. Classification Tree: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

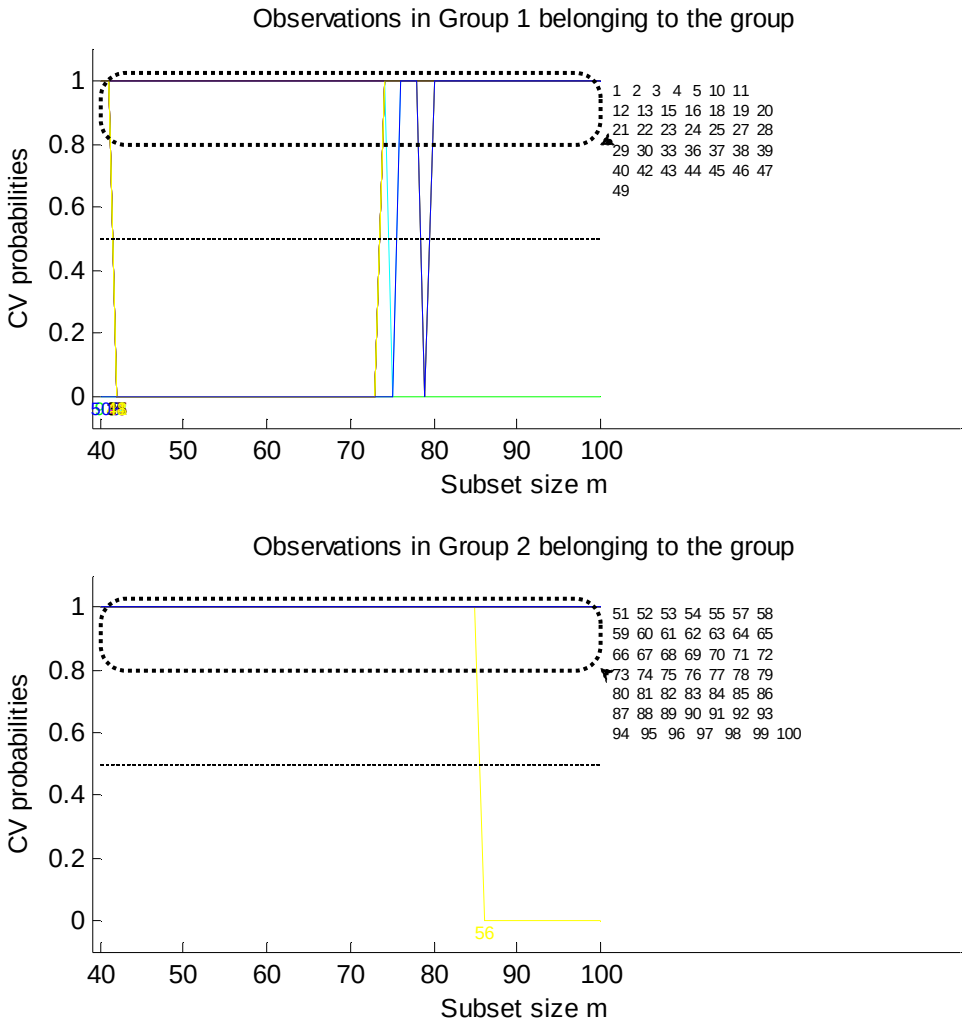

Figure B22. Electrodes data. Classification Tree: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

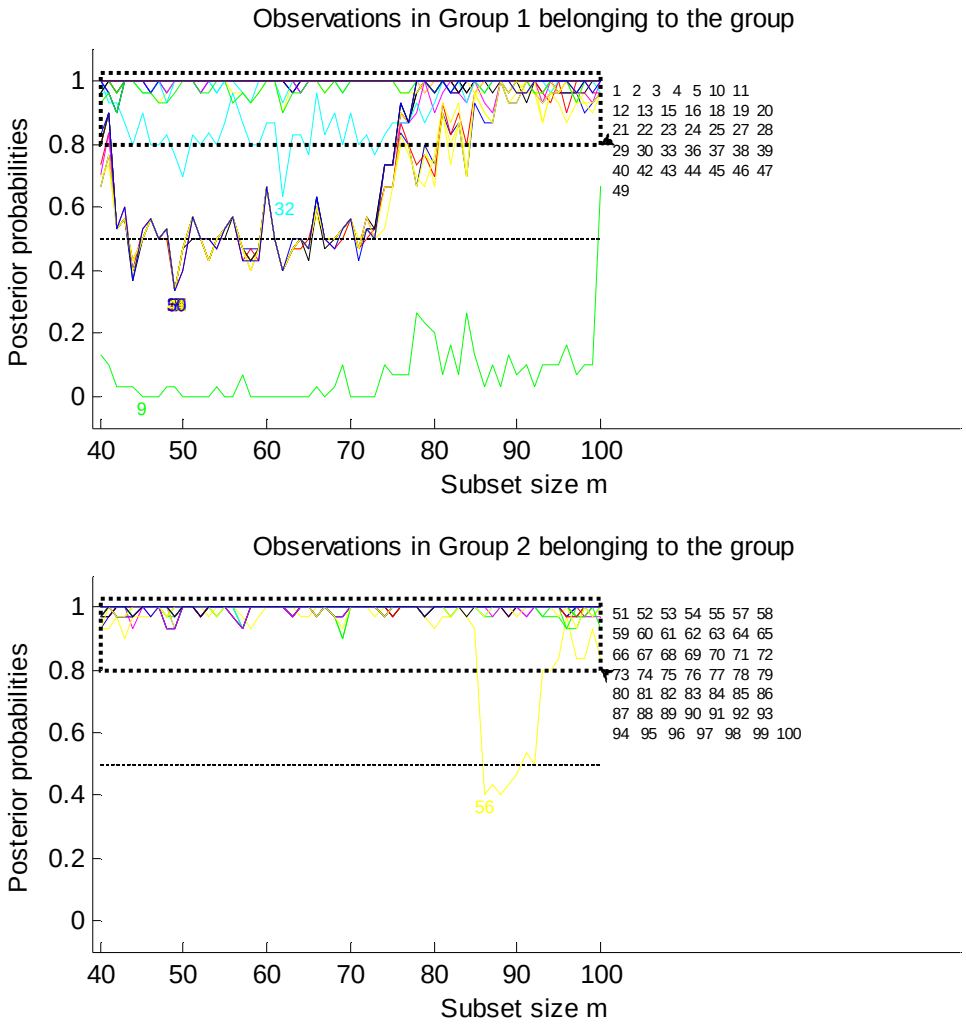

Figure B23. Electrodes data. Bagged Classification Tree: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

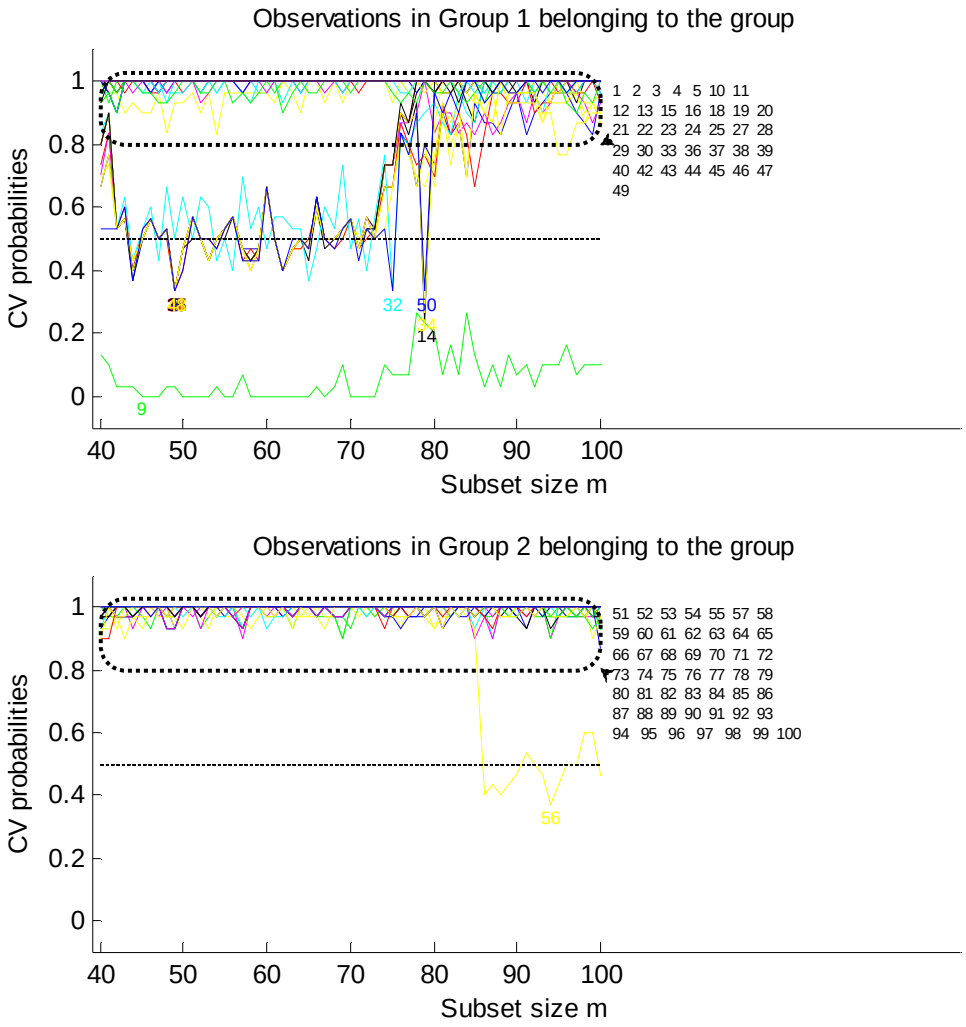

Figure B24. Electrodes data. Bagged Classification Tree: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search

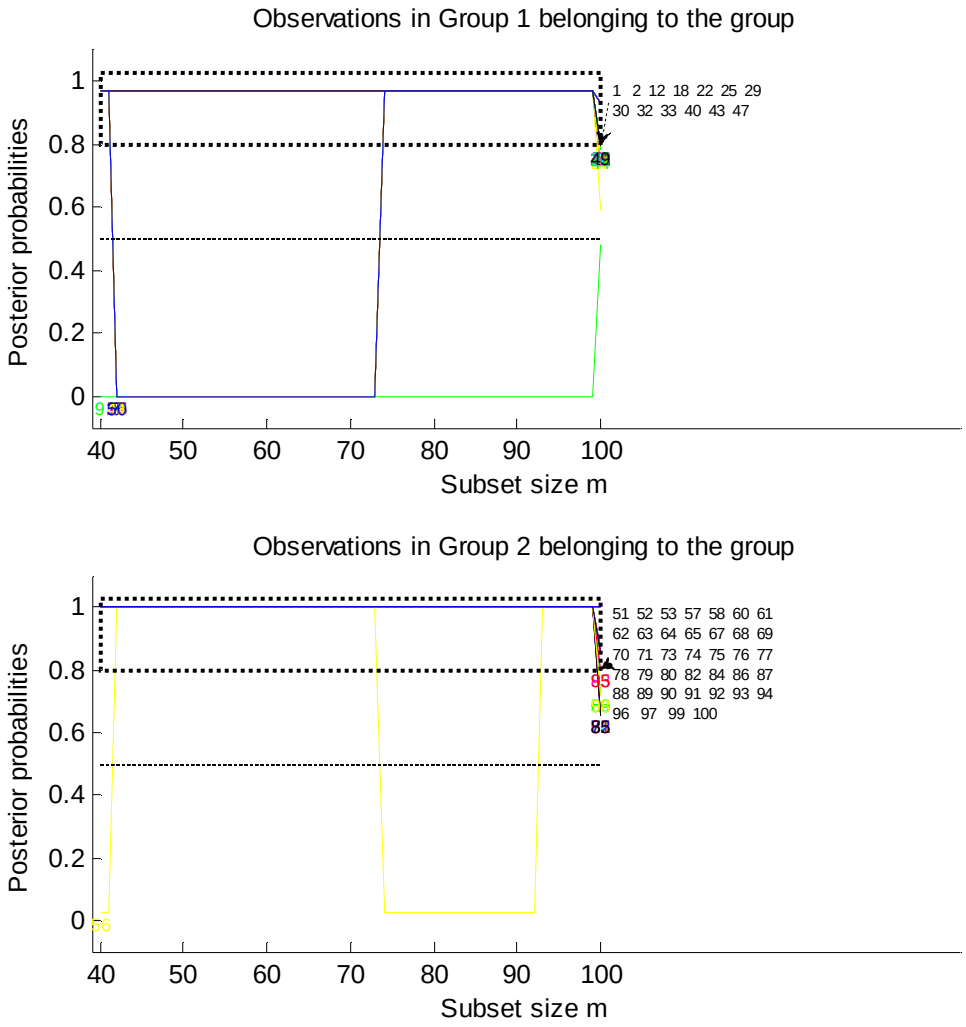

Figure B25. Electrodes data. Adaboost: posterior probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.

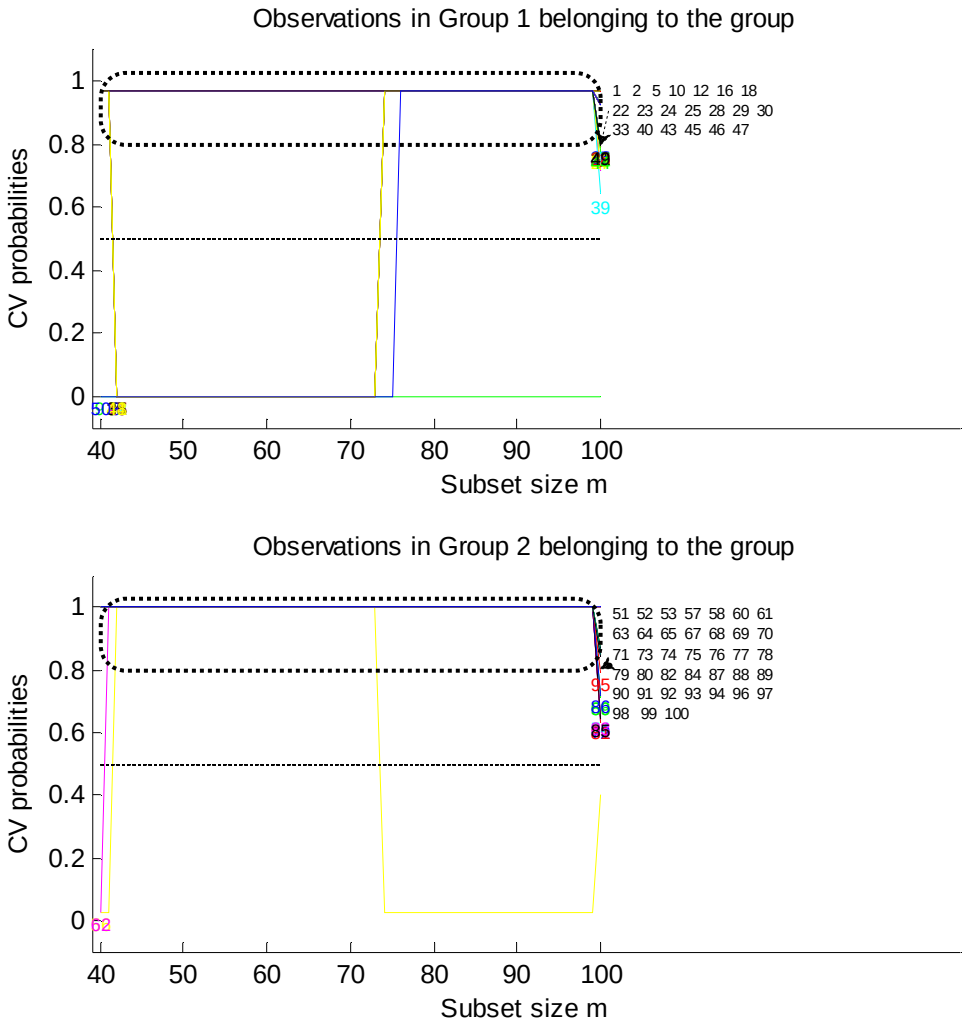

Figure B26. Electrodes data. Adaboost: cross validation probabilities of correct classification of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.

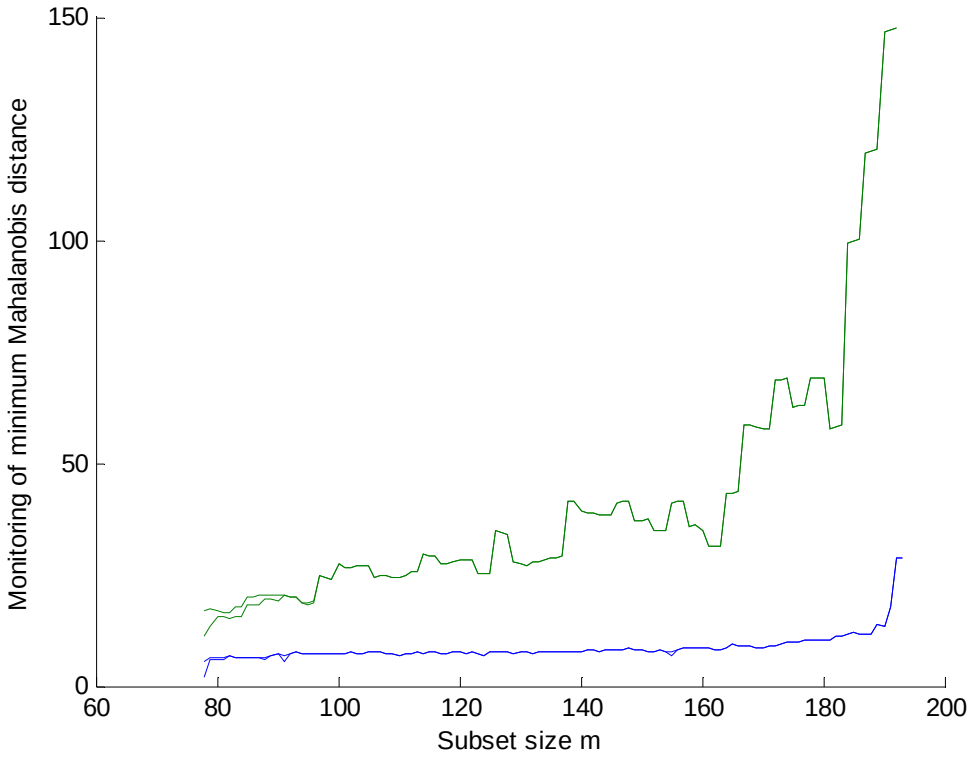

Figure B27. Muscular Dystrophy Data.. Plot of miniimum Mahalanobis distance for units outside the subset. Balanced search

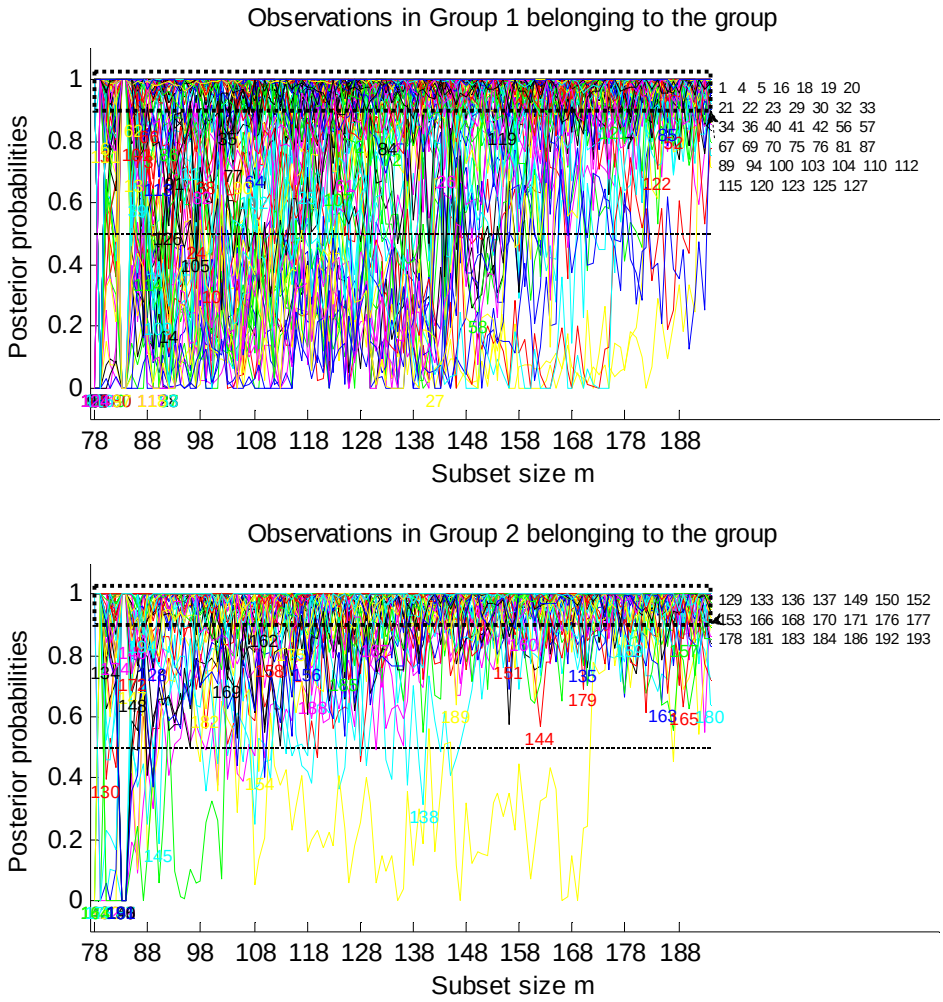

Figure B28. Muscular Dystrophy data. Stacking scheme: posterior probabilities of correct classifcation of the units of Group 1 (upper panel) and of Group 2 (bottom panel). Balanced search.

.

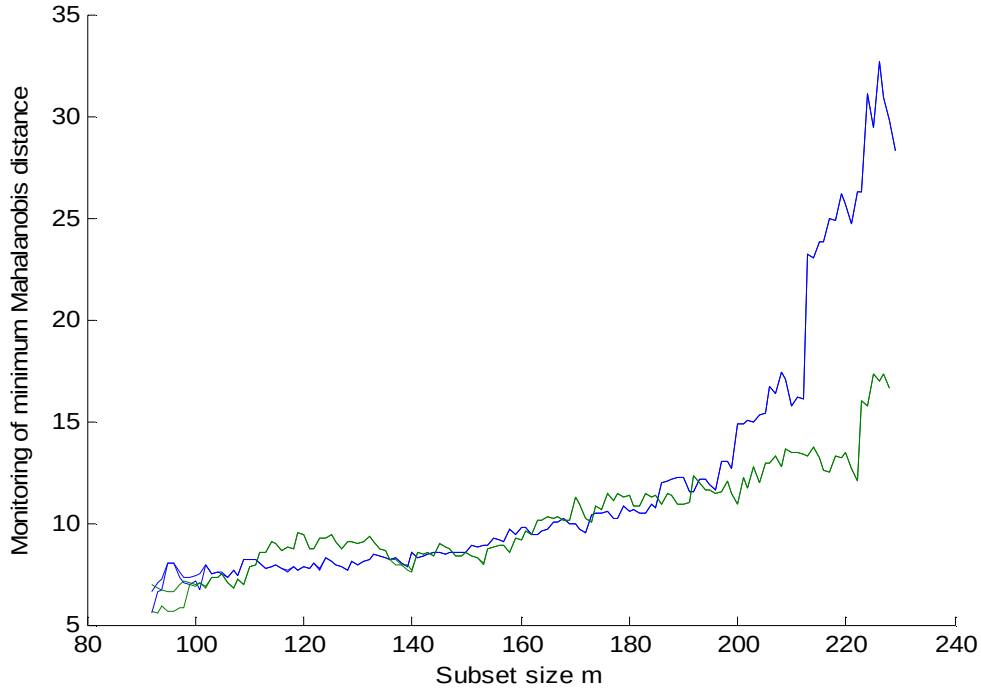

Figure B29. Off-the-book workforce data. Plot of minimum Mahalanobis distances for units outside the subset. Balanced search

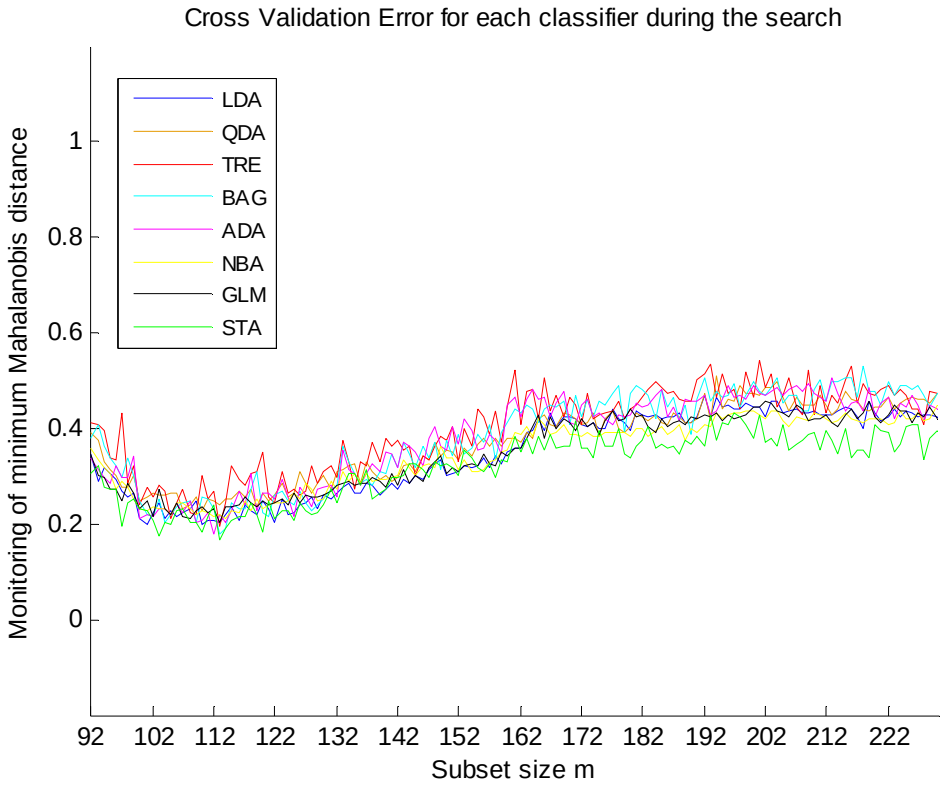

Figure B30. Off-the-book workforce data. Cross validation error for units inside the subset. Balanced search

Bibliography

Atkinson A.C., Riani M. and Cerioli A., (2004). *Exploring Multivariate Data With the Forward Search*. Springer Verlag, New York.

Atkinson A.C, Riani M. and Cerioli A., (2010). The forward search: theory and data analysis (with discussion), *Journal Of The Korean Statistical Society*, vol. 39, p. 117-134.

Boser , B.E., Guyon I., and Vapnik V. (1992), A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, 144-152. ACM Press

Box, G.E. P. and Cox D.R. (1964), An analysis of transformations (with discussion). *Journal of the Royal Statistical Society, Series B 26*, 211-246

Breiman, L., J. Friedman, R. Olshen, and C. Stone, (1984) *Classification and Regression Trees*. Boca Raton, FL: CRC Press.

Breiman, L., (1996), Bagging Predictors. *Machine Learning*, Vol. 24, No. 2, pp. 123-140.

Breiman, L., (1996a), Stacked Regressions. *Machine Learning*, Vol. 24, pp. 49-64.

Breiman, L. (2001). Random forests, *Machine Learning* 45: 5–32.

Breiman, L. and Spector, P. (1992). Submodel selection and evaluation in regression: the X random case, *International Statistical Review* 60: 291–319.

Busa-Fekete R., Kégl B., Elteto t., Szarvas G. (2011), Ranking by calibrated AdaBoost. *JMLR: Workshop and Conference Proceedings 14*, 37-48.

Chang C. and Lin C., (2001), *LIBSVM, a library for support vector machines.* Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.

Cortes C., and Vapnik V. (1995), Support Vector network. *Machine Learning*, 20:273-297.

Demsar J., Statistical Comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7, 1-30.

Dietterich, T. G. (1998), Approximate statistical test for comparing supervised classification learning algorithms. *Neural Computation*, 107, 1895–1923.

Dietterich, T. G. (2000), Ensemble methods in machine learning. In *Proceedings of the First International Workshop on Multiple Classifier Systems* (pp. 1–15). Berlin, Springer.

Dzeroski, S., and Zenko, B. (2002), Is combining classifiers better than selecting the best one? In *Proceedings of the Nineteenth International Conference on Machine Learning*, San Francisco: Morgan Kaufmann.

Dzeroski, S., and Zenko, B. (2002), Stacking with multi-response model trees. In *Multiple Classifiers Systems,Proceedings of the Third International Workshop*, Berlin, Springer

Freund, Y. and Schapire R.E. (1996), *Experiments with a new boosting algorithm, Proceedings of the International Conference on Machine Learning,* pages 148-156, Morgan Kaufmann, San Francisco.

Guyon I., et. al. (2006), Performance prediction challenge. *International Joint Conference on Neural Networks*. Vancouver, Canada

Hastie, T., Tibshirani, R., Friedman, J. (2003), *The Elements of Statistical Learning*. Springer, Heidelberg (2003)

Hastie T., Tibshirani R., and J. Friedman, (2009), *The elements of statistical learning* (2nd ed.). New York, Springer-Verlag.

Hoerl, A. E. and Kennard, R. (1970). Ridge regression: biased estimation for non orthogonal problems, *Technometrics* **12**: 55–67.

Kittler, J., Hatef M., Duin R.P.W. and Matas J. (1998), On combining classifiers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 20, No. 3.

Kohavi R. (1995), A study of cross-validation and Bootstrap for accuracy estimation and model selection. *International Joint Conference on Artificial Intelligence*.

Kuncheva, L.I., Whitaker, C.J., (2003), Measures of Diversity in Classifier Ensembles and their Relationship with the Ensemble Accuracy, *Machine Learning*, 51, 181-207.

Kuncheva, L.I., (2004), *Combining Pattern Classifiers: Methods and Algorithms.* Wiley- Interscience, Hoboken.

LeBlanc, M., and Tibshirani R., (1993), Combining Estimates in Regression and Classification. In *Technical Report 9318*. Department of Statistics, University of Toronto

Ledezma A., Aler R., and Borrajo D., (2001), Empirical study of stacking state- space. *In Thirteenth International Conference on Tools with Artificial Intelligence*, ICTAI'01, pages 210–217.

Mardia, K., Kent, J. and Bibby, J. (1979). *Multivariate Analysis*, Academic Press.

Merz, C. J. (1999), Using correspondence analysis to combine classifiers. *Machine Learning, 36:1/2*, 33–58.

Reid S., and Grudic G. (2009), Regularized linear Models in Stacked Generalization. In Beneditksson J.A., et. al (Eds): *MCS 2009, LNCS 5519*, 112-121

Riani M., and Atkinson A.C., (2001), A unified approach to outliers, influence and transformations in Discriminant Analysis, *Journal of Computational and Graphical Statistics* , Volume 10, Number 3, pages 513-544

Riani, M., Atkinson A.C., and Cerioli A., (2009), Finding an unknown number of multivariate outliers, *Journal of the Royal Statistical Society,* 71, Part 2.

Riani, M., and Zani S. (1997). An iterative method for the detection of multivariate outliers. *Metron 55*, 101-117.

Roli F., Giacinto G., and Vernazza G., (2001), Methods for designing multiple classifier systems. In MCS '01: *Proceedings of the Second International Workshop on Multiple Classifier Systems*, pages 78-87, London, UK, Springer-Verlag.

Schaffer, C., (1993), Selecting a classification method by cross-validation. *Machine Learning 13(1)* (1993) 135–143.

Schaffer C., (1994) Cross-validation, stacking, and bi-level stacking: Meta-methods for classification learning. *In P. Cheeseman & R. W. Odford (Eds.), Selecting models from data: Artificial intelligence and statistics IV*, New York, NY,. Springer-Verlag

Seewald A.K. (2002), Exploring the Parameter State Space of Stacking, in *Proceedings of the 2002 IEEE International Conference on Data Mining (ICDM-2002)*, Maebashi City, Japan.

Seewald A.K. (2002), How to make Stacking better and faster while also taking care of an unknown weakness. In *Proceedings of the 19th International Conference on Machine Learning, ICML-2002*. Morgan Kaufmann Publisher, San Francisco.

Tibshirami R.J., and Tibshirami R., (2009), A Bias correction for the minimum error rate in cross-validation. *The Annals of Applied Statistics*, Vol. 3, No. 2, 822-829.

Ting K. and Witten I., (1997), Stacked generalization: when does it work? In *Proceedings of the International Joint Conference on Artificial Intelligence*.

Ting, K. M. and Witten, I. H. (1999), Issues in stacked generalization. *Journal of Artificial Intelligence Research 10*, pages 271-289.

Todorovski, L., & Dzeroski, S. (2000), Combining multiple models with meta decision trees. In *Proceedings of the Fourth European Conference on Principles of Data Mining and Knowledge Discovery* (pp. 54–64). Berlin, Springer.

Todorovski, L., & Dzeroski, S. (2002), Combining classifiers with meta decision trees. *Machine Learning, 50:3*, 223–249.

Varma S., and Simon R., (2006), Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics*, 91

Witten I.H., Frank E. and Hall M., (2011), *Data Mining. Practical Machine Learning Tools and Techniques*. Third Edition. Morgan Kaufmann

Wolpert, D.H. (1992), Stacked Generalization. *Neural Networks*, Vol. 5, pp. 241-259, Pergamon Press.

Zani, S. (1996), Analisi dei dati statistici. Vol. II. Ossservazioni multidimensionali. Milano, Giuffrè Editore.

Zani, S., Riani. M., and Corbellini A. (1998), Robust bivariate boxplots and multiple outlier detection. *Computational Statistics and Data Analysis 28*, 257-270.

Zenko, B., & Dzeroski, S. (2002), Stacking with an extended set of meta-level attributes and MLR. In *Proceedings of the Thirteenth European Conference on Machine Learning*, Berlin, Springer.

Zenko, B., Todorovski, L., & Dzeroski, S. (2001), A comparison of stacking with MDTs to bagging, boosting, and other stacking methods. In *Proceedings of the First IEEE International Conference on Data Mining* (pp. 669–670). Los Alamitos, IEEE Computer Society