

Luiz S. Oliveira · Marisa Morita · Robert Sabourin

Feature selection for ensembles applied to handwriting recognition

Received: 8 October 2004 / Revised: 19 September 2005 / Accepted 15 November 2005
© Springer-Verlag 2006

Abstract Feature selection for ensembles has shown to be an effective strategy for ensemble creation due to its ability of producing good subsets of features, which make the classifiers of the ensemble disagree on difficult cases. In this paper we present an ensemble feature selection approach based on a hierarchical multi-objective genetic algorithm. The underpinning paradigm is the “overproduce and choose”. The algorithm operates in two levels. Firstly, it performs feature selection in order to generate a set of classifiers and then it chooses the best team of classifiers. In order to show its robustness, the method is evaluated in two different contexts: supervised and unsupervised feature selection. In the former, we have considered the problem of handwritten digit recognition and used three different feature sets and multi-layer perceptron neural networks as classifiers. In the latter, we took into account the problem of handwritten month word recognition and used three different feature sets and hidden Markov models as classifiers. Experiments and comparisons with classical methods, such as Bagging and Boosting, demonstrated that the proposed methodology brings compelling improvements when classifiers have to work with very low error rates. Comparisons have been done by considering the recognition rates only.

Keywords Ensemble of classifiers · Feature selection · Handwriting recognition · Multi-objective optimization · Genetic algorithms

L. S. Oliveira (✉)
Pontifícia Universidade Católica do Paraná (PUCPR), Programa de Pós-Graduação em Informática Aplicada (PPGIA), Rua Imaculada Conceição 1155, Prado Velho, 80215-901, Curitiba, Pr, Brazil
E-mail: soares@ppgia.pucpr.br

M. Morita · R. Sabourin
École de Technologie Supérieure (ETS), Laboratoire d’Imagerie, de Vision et d’Intelligence Artificielle 1100, rue Notre Dame Ouest, Montreal, Canada, H3C 1K3

Present address:

M. Morita
HSBC Bank Brazil, IT, Curitiba, PR, Brazil
E-mail: marisa.e.morita@hsbc.com.br

1 Introduction

Ensemble of classifiers has been widely used to reduce model uncertainty and improve generalization performance. Developing techniques for generating candidate ensemble members is a very important direction of ensemble of classifiers research. It has been demonstrated that a good ensemble is one where the individual classifiers in the ensemble are both accurate and make their errors on different parts of the input space (there is no gain in combining identical classifiers) [1–4]. In other words, an ideal ensemble consists of good classifiers (not necessarily excellent) that disagree as much as possible on difficult cases.

The literature has shown that varying the feature subsets used by each member of the ensemble should help to promote this necessary diversity [4–7]. Traditional feature selection algorithms aim at finding the best trade-off between features and generalization. On the other hand, ensemble feature selection has the additional goal of finding a set of feature sets that will promote disagreement among the component members of the ensemble. The Random Subspace Method (RMS) proposed by Ho in [5] was one early algorithm that constructs an ensemble by varying the subset of features. More recently some strategies based on genetic algorithms (GAs) have been proposed [4, 8, 9]. All these strategies claim better results than those produced by traditional methods for creating ensembles such as Bagging and Boosting. In spite of the good results brought by GA-based methods, they still can be improved in some aspects, e.g., avoiding classical methods such as the weighted sum to combine multiple objective functions. It is well known that when dealing with this kind of combination, one should deal with problems such as scaling and sensitivity towards the weights.

It has been demonstrated that feature selection through multi-objective genetic algorithm (MOGA) is a very powerful tool for finding a set of good classifiers [10, 11], since GA is quite effective in rapid global search of large, non-linear and poorly understood spaces [12]. Besides, it can overcome problems such as scaling and sensitivity towards

the weights. Kudo and Sklansky [13] have compared several algorithms for feature selection and concluded that GAs are suitable when dealing with large-scale feature selection (number of features is over 50). This is the case of most of the problems in handwriting recognition, which is the test problem in this work.

In this light, we propose an ensemble feature selection approach based on a hierarchical MOGA. The underlying paradigm is “overproduce and choose” [14, 15]. The algorithm operates in two levels. The former is devoted to generate a set of good classifiers by minimizing two criteria: error rate and number of features. The latter combines these classifiers in order to find an ensemble by maximizing the following two criteria: accuracy of the ensemble and a measure of diversity.

Recently, the issue of using diversity to build ensemble of classifiers has been widely discussed. Several works have demonstrated that there is a weak correlation between diversity and ensemble performance [16, 17]. In light of this, some authors have claimed that diversity brings no benefits in building ensemble of classifiers [18], on the other hand, others suggest that the study of diversity in classifier combination might be one of the lines for further exploration [19].

In spite of the weak correlation between diversity and performance, we argue that diversity might be useful to build ensembles of classifiers. We demonstrated through experimentation that using diversity jointly with performance to guide selection can avoid overfitting during the search. In order to show robustness of the proposed methodology, it was evaluated in two different contexts: supervised and unsupervised feature selection. In the former, we have considered the problem of handwritten digit recognition and used three different feature sets and multi-layer perceptron (MLP) neural networks as classifiers. In such a case, the classification accuracy is supplied by the MLPs in conjunction with the sensitivity analysis. This approach makes it feasible to deal with huge databases in order to better represent the pattern recognition problem during the fitness evaluation. In the latter, we took into account the problem of handwritten month word recognition and used three different feature sets and hidden Markov models (HMM) as classifiers. We demonstrate that it is feasible to find compact clusters and complementary high-level representations (codebooks) in subspaces without using the recognition results of the system. Experiments and comparisons with classical methods, such as Bagging and Boosting, demonstrated that the proposed methodology brings compelling improvements when classifiers have to work with very low error rates. It is worth of remark that the comparisons have been done by considering the recognition rates only.

The remainder of this paper is organized as follows. Section 2 presents a brief review about the methods for ensemble creation. Section 3 provides an overview of the strategy. Section 4 introduces briefly the multi-objective genetic algorithm we are using in this work. Section 5 describes the classifiers and feature sets for both supervised and unsupervised contexts. Section 6 introduces how we

have implemented both levels of the proposed methodology and Sect. 7 reports the experimental results. Finally, Sect. 8 discusses the reported results and Sect. 9 concludes the paper.

2 Related works

Assuming the architecture of the ensemble as the main criterion, we can distinguish among serial, parallel, and hierarchical schemes, and if the classifiers of the ensemble are selected or not by the ensemble algorithm we can divide them into selection-oriented and combiner-oriented methods [20, 21]. Here we are more interested in the first class, which try to improve the overall accuracy of the ensemble by directly boosting the accuracy and the diversity of the experts of the ensemble. Basically, they can be divided into resampling methods and feature selection methods.

Resampling techniques can be used to generate different hypotheses. For instance, bootstrapping techniques [22] may be used to generate different training sets and a learning algorithm can be applied to the obtained subsets of data in order to produce multiple hypotheses. These techniques are effective especially with unstable learning algorithms, which are algorithms very sensitive to small changes in the training data. In bagging [23] the ensemble is formed by making bootstrap replicates of the training sets, and then multiple generated hypotheses are used to get an aggregated predictor. The aggregation can be performed by averaging the outputs in regression or by majority or weighted voting in classification problems.

While in bagging the samples are drawn with replacement using a uniform probability distribution, in boosting methods [24] the learning algorithm is called at each iteration using a different distribution or weighting over the training examples. This technique places the highest weight on the examples most often misclassified by the previous base learner: in this manner the classifiers of the ensemble focus their attention on the hardest examples. Then the boosting algorithm combines the base rules taking a weighted majority vote of the base rules.

The second class of methods regards those strategies based on feature selection. The concept behind these approaches consists of reducing the number of input features of the classifiers, a simple method to fight the effects of the classical curse of dimensionality problem. For instance, the random subspace method [5, 7] relies on a pseudorandom procedure to select a small number of dimensions from a given feature space. In each pass, such a selection is made and a subspace is fixed. All samples are projected to this subspace, and a classifier is constructed using the projected training samples. In the classification a sample of an unknown class is projected to the same subspace and classified using the corresponding classifier. In the same vein of the random subspace method lies the input decimation method [25], which reduces the correlation among the errors of the base classifiers, by decoupling the classifiers by

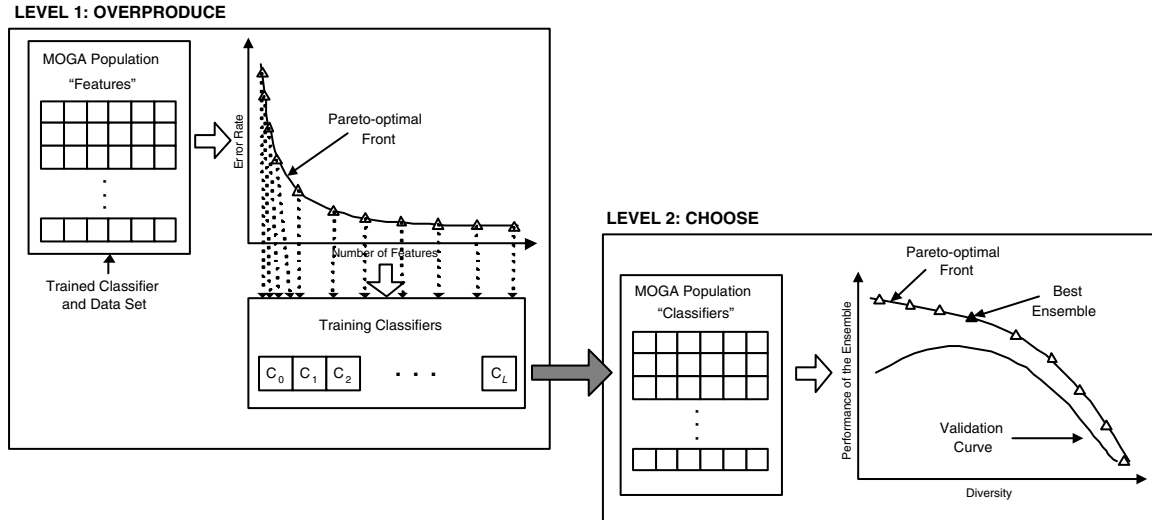


Fig. 1 An overview of the proposed methodology

training them with different subsets of the input features. It differs from the random subspace as for each class the correlation between each feature and the output of the class is explicitly computed, and the classifier is trained only on the most correlated subset of features.

Recently, several authors have investigated GA to design ensemble of classifiers. Kuncheva and Jain [26] suggest two simple ways to use genetic algorithm to design an ensemble of classifiers. They present two versions of their algorithm. The former uses just disjoint feature subsets while the latter considers (possibly) overlapping feature subsets. The fitness function employed is the accuracy of the ensemble, however, no measure of diversity is considered. Gerra-Salcedo and Withley [8] used a simple GA to explore the space of all possible feature subsets, and then create an ensemble based on them. In their experiments, this GA-based approach outperformed classical methods such as Bagging and Boosting. In spite of the fact they achieved interesting results, they did not consider any measure of diversity. A more elaborate method, also based on GA, was proposed by Optiz [4]. In his work, he stresses the importance of a diversity measure by including it in the fitness calculation. The drawback of this method is that the objective functions are combined through the weighted sum. It is well known that when dealing with this kind of combination, one should deal with problems such as scaling and sensitivity towards the weights. More recently Gunter and Bunke [27] have applied feature selection in conjunction with floating search to create ensembles of classifiers for the field of handwriting recognition. They used handwritten words and HMMs as classifiers to evaluate their algorithm. The feature set was composed of nine discrete features, which makes simpler the feature selection process. A drawback of this method is that one must set a priori the number of classifiers in the ensemble.

3 Methodology overview

In this section, we outline the hierarchical approach proposed. As stated before, it is based on an “overproduce and choose” paradigm where the first level generates several classifiers by conducting feature selection and the second one chooses the best ensemble among such classifiers. Figure 1 depicts the proposed methodology. Firstly, we carry out feature selection by using a MOGA. It gets as inputs a trained classifier and its respective data set. Since the algorithm aims at minimizing two criteria during the search,¹ it will produce at the end a two-dimensional Pareto-optimal front, which contains a set of classifiers (trade-offs between the criteria being optimized). The final step of this first level consists in training such classifiers.

Once the set of classifiers have been trained, the second level is suggested to pick the members of the team which are most diverse and accurate. Let $A = C_1, C_2, \dots, C_L$ be a set of L classifiers extracted from the Pareto-optimal and B a chromosome of size L of the population. The relationship between A and B is straightforward, i.e., the gene i of the chromosome B is represented by the classifier C_i from A . Thus, if a chromosome has all bits selected, all classifiers of A will be included in the ensemble. Therefore, the algorithm will produce a 2-dimensional Pareto-optimal front which is composed of several ensembles (trade-offs between accuracy and diversity). In order to choose the best one, we use a validation set, which points out the most diverse and accurate team among all. Later in this paper, we will discuss the issue of using diversity to choose the best ensemble.

In both cases, MOGAs are based on bit representation, one-point crossover, and bit-flip mutation. In our experiments, MOGA used is the Non-dominated Sorting Genetic

¹ Error rate and number of features in the case of supervised feature selection and a clustering index and the number of features in the case of unsupervised feature selection.

Algorithm (NSGA) with elitism proposed by Srinivas and Deb in [29], which is briefly introduced in the next section.

4 Multi-objective genetic algorithm

Since the concept of multi-objective genetic algorithm (MOGA) will be explored in the remaining of this paper, this section briefly introduces it.

A general multi-objective optimization problem consists of a number of objectives and is associated with a number of inequality and equality constraints. Solutions to a multi-objective optimization problem can be expressed mathematically in terms of nondominated points, i.e., a solution is dominant over another only if it has superior performance in all criteria. A solution is said to be Pareto-optimal if it cannot be dominated by any other solution available in the search space. In our experiments, the algorithm adopted is the Nondominated Sorting Genetic Algorithm (NSGA) with elitism proposed by Srinivas and Deb in [28, 29].

The idea behind the NSGA is that a ranking selection method is used to emphasize good points and a niche method is used to maintain stable subpopulations of good points. It differs from simple GA only in the way the selection operator works. The crossover and mutation remain as usual. Before the selection is performed, the population is ranked based on an individual's nondomination. The nondominated individuals present in the population are first identified from the current population. Then, all these individuals are assumed to constitute the first nondominated front in the population and assigned a large dummy fitness value. The same fitness value is assigned to give an equal reproductive potential to all these nondominated individuals. This is exemplified in Fig. 2a. In such a case, a population of six individuals was classified into three nondominated fronts and each individual of the first front received a large dummy fitness (6.00 in this example).

In order to maintain the diversity in the population, these classified individuals are then shared with their dummy fitness values. Sharing is achieved by performing selection operation using degraded fitness values obtained

by dividing the original fitness value of an individual by a quantity proportional to the number of individuals around it. After sharing, these nondominated individuals are ignored temporarily to process the remaining population in the same way to identify individuals for the second nondominated front. These new sets of points are then assigned a new dummy fitness which is kept smaller than the minimum shared dummy fitness of the previous front. This process is continued until the entire population is classified into several fronts. This process is illustrate in Fig. 2b. It can be observed from this Figure that the individuals "1" and "3" had their fitness shared because they are close to each other. In this case, their fitness were reduced from 6.00 to 4.22. Then, the dummy fitness is assigned to the individuals of the second front by multiplying the lowest value of the first front by a constant k (let us say $k = 0.95$ for this example). Therefore, the individuals of the second front will receive a dummy fitness of 4.00 (4.22×0.95). Since the two individuals of the second front are not close to each other, their dummy fitness is maintained and a dummy fitness is assigned to the individual of the last front ($4.00 \times 0.95 = 3.80$).

The population is then reproduced according to the dummy fitness values. Since individuals in the first front have the maximum fitness value, they get more copies than the rest of the population. This was intended to search for the nondominated regions of Pareto-optimal fronts. The efficiency of NSGA lies in the way multiple objectives are reduced to a dummy fitness function using nondominated sorting procedures.

Figure 3 shows a flow chart of NSGA. The algorithm is similar to a simple GA except for the classification of nondominated fronts and the sharing operation. The sharing in each front is achieved by calculating a sharing function value between two individuals in the same front as:

$$Sh(d(i, j)) = \begin{cases} 1 - \left(\frac{d(i, j)}{\sigma_{share}}\right)^2 & \text{if } d(i, j) < \sigma_{share} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where $d(i, j)$ is the distance between two individuals i and j in the current front and σ_{share} is the maximum distance allowed between any two individuals to become members

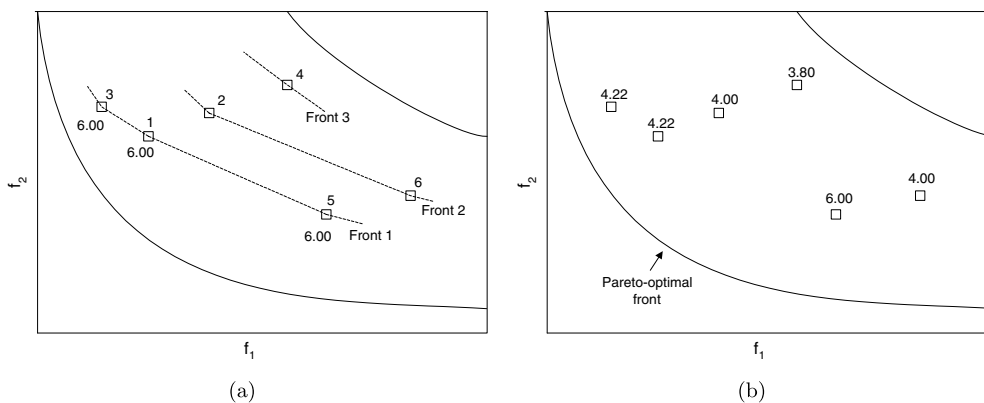


Fig. 2 Sorting population: **a** The population is classified into three nondominated fronts and **b** shared fitness values of six solutions

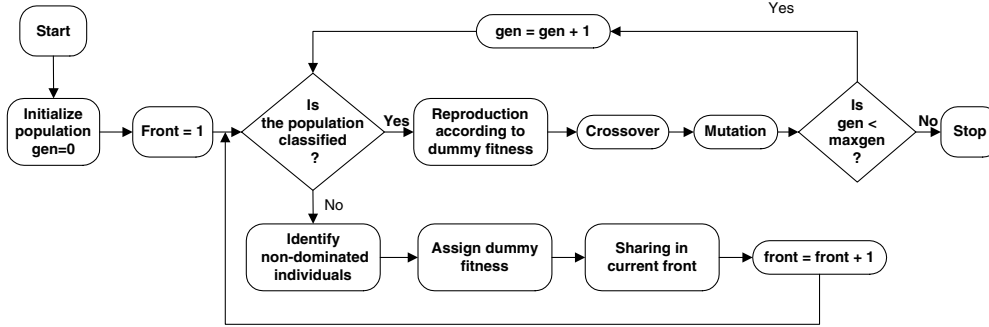


Fig. 3 Flow chart of NSGA

of a niche. In any application of sharing, we can implement either genotypic sharing, since we always have a genotype (the encoding), or phenotypic sharing. However, Deb and Goldberg in [30] indicate that in general, phenotypic sharing is superior to genotypic sharing. Thus, we have used a phenotypic sharing which is calculated from the normalized Euclidean distance between the objective functions.

The parameter σ_{share} can be calculated as follows [30]:

$$\sigma_{\text{share}} \approx \frac{0.5}{\sqrt[q]{q}} \quad (2)$$

where q is the desired number of distinct Pareto-optimal solutions and p is the number of decision variables. Although the calculation of σ_{share} depends on this parameter q , it has been shown [29] that the use of the above equation with $q \approx 10$ works in many test problems.

5 Classifiers and feature sets

As stated before, we have carried out experiments in both supervised and unsupervised contexts. The remaining of this section describes the feature sets and classifiers we have used.

5.1 Supervised context

To evaluate the proposed methodology in the supervised context, we have used three base classifiers trained to recognize handwritten digits of NIST SD19. Such classifiers were trained with three well-known feature sets: Concavities and Contour (CC_{sc}) [31], Distances (DDD_{sc}) [32], and Edge Maps (EM_{sc}) [33]. It should be noted, though, that the original feature set of distances proposed by Oh and Suen [32] contains 256 features. After carrying out some experiments with different strategies of zoning, we realized that using 96 features (6 zones: 3 horizontal and 2 vertical) we could achieve the same results as using 256 features (16 symmetrical zones).

All classifiers here are MLPs trained with the gradient descent applied to a sum-of-squares error function [34].

The transfer function employed is the familiar sigmoid function. In order to monitor the generalization performance during learning and terminate the algorithm when there is no longer an improvement, we have used the method of cross-validation. Such a method takes into account a validation set, which is not used for learning, to measure the generalization performance of the network. During learning, the performance of the network on the training set will continue to improve, but its performance on the validation set will only improve to a point, where the network starts to overfit the training set, that the learning algorithm is terminated. All networks have one hidden layer where the units of input and output are fully connected with units of the hidden layer, where the number of hidden units were determined empirically (see Table 1). The learning rate and the momentum term were set at high values in the beginning to make the weights quickly fit the long ravines in the weight space, then these parameters were reduced several times according to the number of iterations to make the weights fit the sharp curvatures.

Among the different strategies of rejection we have tested, the one proposed by Fumera et al [35] provided the better error-reject trade-off for our experiments. Basically, this technique suggests the use of multiple reject thresholds for the different data classes (T_0, \dots, T_n) to obtain the optimal decision and reject regions. In order to define such thresholds we have developed an iterative algorithm, which takes into account a decreasing function of the threshold variables $R(T_0, \dots, T_n)$ and a fixed error rate T_{error} . We start from all threshold values equal to 1, i.e., the error rate equal to zero since all images are rejected. Then, at each step, the algorithm decreases the value of one of the thresholds in order to increase the accuracy until the error rate exceeds T_{error} .

Table 1 Description and performance of the classifiers on TSDB (zero-rejection level)

Feature set	Number of features	Units in the hidden layer	Rec. rate (%)
CC _{sc}	132	80	99.13
DDD _{sc}	96	60	98.17
EM _{sc}	125	70	97.04




Fig. 4 Performance of the classifiers on the test set for error rates varying from 0.10 to 0.50%. **a** Trade-off between recognition rate and error rate. **b** Trade-off between reject rate and error rate

The training (TRDB_{sc}) and validation (VLDB1_{sc}) sets are composed of 195,000 and 28,000 samples from hsf_0123 series respectively while the test set (TSDB_{sc}) is composed of 30,089 samples from the hsf_7. We consider also a second validation set (VLDB2_{sc}), which is composed of 30,000 samples of hsf_7. This data is used to select the best ensemble of classifiers. Figure 4 shows the performance on the test set of all classifiers for error rates varying from 0.10 to 0.50%, while Table 1 reports the performance of all classifiers at zero-rejection level. The curves depicted in Fig. 4 are much more meaningful when dealing with real applications since they describe the recognition rate in relation to a specific error rate, including implicitly a corresponding reject rate. This rate also allows us to compute the reliability of the system for a given error rate. It can be done by using Eq. (3).

$$\text{Reliability} = \frac{\text{Rec.rate}}{\text{Rec.rate} + \text{Error rate}} \times 100 \quad (3)$$

Figure 4 corroborates that recognition of handwritten digits is still an open problem when very low error rates are required. Consider for example our best classifier, which reaches 99.13% at zero-rejection level on the test set. If we allow an error rate of 0.1%, i.e., just one error in 1,000, the recognition rate of such classifier drops from 99.13% to 91.83%. This means that we have to reject 8.07% to get 0.1% of error (Fig. 4a). We will demonstrate that the ensemble of classifiers can significantly improve the performance of the classifiers for low error rates.

5.2 Unsupervised context

To evaluate the proposed methodology in unsupervised context we have used three HMM-based classifiers trained to recognize handwritten Brazilian month words (“Janeiro”, “Fevereiro”, “Maro”, “Abril”, “Maio”, “Junho”, “Julho”, “Agosto”, “Setembro”, “Outubro”, “Novembro”, “Dezembro”). The training (TRDB_{uc}), validation (VLDB1_{uc}), and

testing (TSDB_{uc}) sets are composed of 1,200, 400, and 400 samples, respectively. In order to increase the training and validation sets, we have also considered 8,300 and 1,900 word images, respectively, extracted from the legal amount database. This is possible because we are considering character models. We consider also a second validation set (VLDB2_{uc}) of 500 handwritten Brazilian month words [36]. Such data is used to select the best ensemble of classifiers.


Given a discrete HMM-based approach, each word image is transformed as a whole into a sequence of observations by the successive application of preprocessing, segmentation, and feature extraction. Preprocessing consists of correcting the average character slant. The segmentation algorithm uses the upper contour minima and some heuristics to split the date image into a sequence of segments (graphemes), each of which consists of a correctly segmented, an under-segmented, or an over-segmented character. A detailed description of the preprocessing and segmentation stages is given in [37].

The word models are formed by the concatenation of appropriate elementary HMMs, which are built at letter and space levels. Two topologies of letter models were chosen based on the output of our grapheme-based segmentation algorithm which may produce a correct segmentation of a letter, a letter under-segmentation or a letter over-segmentation into two, three, or four graphemes depending on each letter. Considering uppercase and lowercase letters, we need 42 models since the legal amount alphabet is reduced to 21 letter classes and we are not considering the unused ones. Thus, regarding the two topologies, we have 84 HMMs which are trained using the Baum-Welch algorithm with the Cross-Validation procedure.

The feature set that feeds the first classifier is a mixture of concavity and contour features (CC_{uc}) [31]. In this case, each grapheme is divided into two equal zones (horizontal) where for each region a concavity and contour feature vector of 17 components is extracted. Therefore, the final feature vector has 34 components. The other two classifiers make


Table 2 Performance of the classifiers on the test set

Feature set	Number of features	Codebook size	Rec. rate (%)
CC _{uc}	34	80	86.1
DDD32 _{uc}	32	40	73.0
DDD64 _{uc}	64	60	64.5


Fig. 5 Zoning based on the reference baselines: **a** baselines and **b** 4-region zoning

use of a feature set based on distances [32]. The former uses the same zoning discussed before (two equal zones), but in this case, for each region a vector of 16 components is extracted. This leads to a final feature vector of 32 components (DDD32_{uc}). For the latter we have tried a different zoning. The grapheme is divided into four zones using the reference baselines (see Fig. 5), hence, we have a final feature vector composed of 64 components (DDD64_{uc}). Table 2 reports the performance of all classifiers on the test set at zero-rejection level. Figure 6 shows the performance of all classifiers for error rates varying from 1% to 4%. The strategy for rejection used in this case is the one discussed previously. We have chosen higher error rates in this case due to the size of the database we are dealing with.

It can be observed from Fig. 6 that the recognition rates with error fixed at 1% are very poor, hence, the number of rejected patterns is very high. We will see in the next sections that the proposed methodology can improve these results considerably.


Fig. 6 Performance of the classifiers on the test set for error rates varying from 1 to 4%. **a** Trade-off between recognition rate and error rate. **b** Trade-off between reject rate and error rate

6 Implementation

This section introduces how we have implemented both levels of the proposed methodology. First we discuss the supervised context and then the unsupervised.

6.1 Supervised context

6.1.1 Supervised feature subset selection

The feature selection algorithm used in here was introduced in [11]. To make this paper self-contained, a brief description is included in this section.

Regarding feature selection algorithms, they can be classified into two categories based on whether or not feature selection is performed independently of the learning algorithm used to construct the classifier. If feature selection is done independently of the learning algorithm, the technique is said to follow a filter approach. Otherwise, it is said to follow a wrapper approach [38]. While the filter approach is generally computationally more efficient than the wrapper approach, its major drawback is that an optimal selection of features may not be independent of the inductive and representational biases of the learning algorithm that is used to construct the classifier. On the other hand, the wrapper approach involves the computational overhead of evaluating candidate feature subsets by executing a given learning algorithm on the database using each feature subset under consideration.

As stated elsewhere, the idea of using feature selection is to promote diversity among the classifiers. To tackle such a task we have to optimize two objective functions: minimization of the number of features and minimization of the error rate of the classifier. Computing the first one is simple, i.e., the number of selected features. The problem lies in computing the second one, i.e., the error rate supplied by the classifier. Regarding a wrapper approach, in each generation,

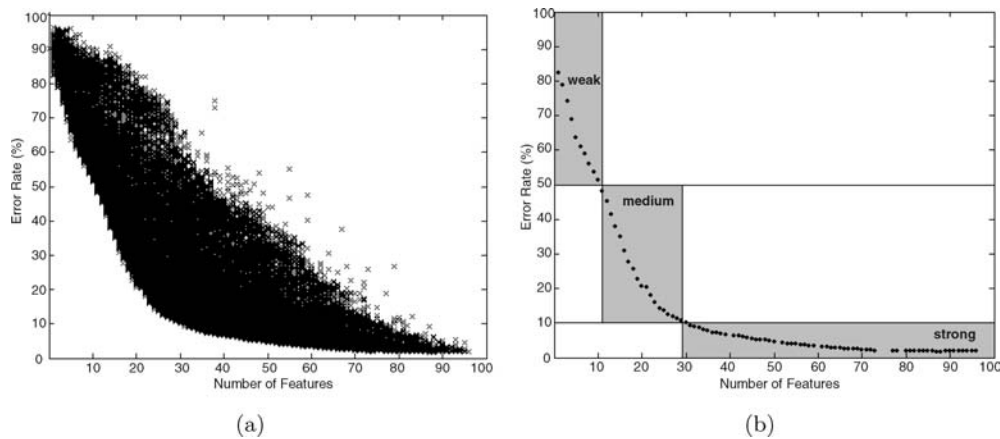


Fig. 7 Supervised feature selection using a Pareto-based approach. **a** Evolution of the population in the objective plane, **b** Pareto-optimal front and its different classes of classifiers

evaluation of a chromosome (a feature subset) requires training the corresponding neural network and computing its accuracy. This evaluation has to be performed for each of the chromosomes in the population. Since such a strategy is not feasible due to the limits imposed by the learning time of the huge training set considered in this work, we have adopted the strategy proposed by Moody and Utans in [39], who use the sensitivity of the network to estimate the relationship between the input features and the network performance.

The sensitivity of the network model to variable β is defined as:

$$S_{\beta} = \frac{1}{N} \sum_{j=1}^N \text{ASE}(\bar{x}_{\beta}) - \text{ASE}(x_{\beta}) \quad (4)$$

with

$$\bar{x}_{\beta} = \frac{1}{N} \sum_{j=1}^N x_{\beta_j} \quad (5)$$

where x_{β_j} is the β th input variable of the j th exemplar. S_{β} measures the effect on the training ASE (average square error) of replacing the β th input x_{β} by its average \bar{x}_{β} . Moody and Utans show that when variables with small sensitivity values with respect to the network outputs are removed, they do not influence the final classification. So, in order to evaluate a given feature subset we replace the unselected features by their averages. In this way, we avoid training the neural network and hence turn the wrapper approach feasible for our problem. We call this strategy modified-wrapper. Such a scheme has been employed also by Yuan et al in [40], and it makes it feasible to deal with huge databases in order to better represent the pattern recognition problem during the fitness evaluation.² Moreover it can accommodate multiple criteria such as the number of features and the accuracy of the classifier, and generate the Pareto-optimal front in the

² If small databases are considered, then a full-wrapper could replace the proposed modified-wrapper.

first run of the algorithm. Figure 7 shows the evolution of the population in the objective plane and its respective Pareto-optimal front.

It can be observed in Fig. 7b that the Pareto-optimal front is composed of several different classifiers. In order to get a better insight about them, they were classified into three different groups: weak, medium, and strong. It can be observed that among all those classifiers there are very good ones. To find out which classifiers of the Pareto-optimal front compose the best ensemble, we carried out a second level of search. Once we did not train the models during the search (the training step is replaced by the sensitivity analysis), the last step of feature selection consists of training the solutions provided by the Pareto-optimal front (1).

Before discussing the second level of our algorithm, we would like to stress the importance of using MOGA for feature selection. Therefore, we show a result achieved by a single GA, where the objective (number of features and error rate of the classifiers) were combined through the weighted-sum approach (Fig. 8). As expected, the results achieved by the weighted-sum approach presented a premature convergence to a specific region of the search space instead of maintaining a diverse population. This kind of behavior can be explained by the sensitivity towards weight presented by the weighted-sum approach. Since we have chosen weights to favor solutions with a small error rate rather than a small number of features, the selection pressure drove the search to the region where the error rates are lower. Thus, after several trials of using different weights we did not succeed in finding the Pareto-optimal front but rather an approximation of the Pareto-optimal solutions.

6.1.2 Choosing the best ensemble

As defined in Sect. 3 each gene of the chromosome is represented by a classifier produced in the previous level. Therefore, if a chromosome has all bits selected, all classifiers of will compose the team. In order to find the best ensemble of classifiers, i.e., the most diverse set of classifiers that brings




Fig. 8 Feature selection using the weighted-sum approach. **a** Evolution of the population in the objectives plane (one trial) and **b** Pareto-optimal solutions found by the classical approach after several trials

a good generalization, we have used two objective functions during this level of the search, namely, maximization of the recognition rate of the ensemble and maximization of a measure of diversity. We have tried different measures such as overlap, entropy [41], and ambiguity [3]. The results achieved with ambiguity and entropy were very similar. In this work we have used ambiguity as diversity measure. The ambiguity is defined as follows:

$$a_i(x_k) = [V_i(x_k) - \bar{V}(x_k)]^2 \quad (6)$$

where a_i is the ambiguity of the i th classifier on the example x_k , randomly drawn from an unknown distribution, while V_i and \bar{V} are the i th classifier and the ensemble predictions, respectively. In other words, it is simply the variance of ensemble around the mean, and it measures the disagreement among the classifiers on input x . Thus the contribution to diversity of an ensemble member i as measured on a set of M samples is:

$$A_i = \frac{1}{M} \sum_{k=1}^M a_i(x_k) \quad (7)$$

and the ambiguity of the ensemble is

$$\bar{A} = \frac{1}{N} \sum A_i \quad (8)$$

where N is the number of classifiers. So, if the classifiers implement the same functions, the ambiguity \bar{A} will be low, otherwise it will be high. In this scenario the error from the ensemble is

$$E = \bar{E} - \bar{A} \quad (9)$$

where \bar{E} is the average errors of the single classifiers and \bar{A} is the ambiguity of the ensemble. Equation (9) expresses the trade-off between bias and variance in the ensemble, but in a different way than the common bias-variance relation in which the averages are over possible training sets instead

of ensemble averages. If the ensemble is strongly biased the ambiguity will be small, because the classifiers implement very similar functions and thus agree in inputs even outside the training set [3].

At this level of the strategy we want to maximize the generalization of the ensemble, therefore, it will be necessary to use a way of combining the outputs of all classifiers to get a final decision. To do this, we have used the average, which is a simple and effective scheme of combining predictions of the neural networks [42]. Other combination rules such as product, min, and max have been tested but the simple average has produced slightly better results. In order to evaluate the objective functions during the search described above we have used the validation set VLDB1sc.

6.2 Unsupervised context

6.2.1 Unsupervised feature subset selection

A lot of work done in the field of handwritten word recognition take into account discrete HMMs as classifiers, which have to be fed with a sequence of discrete values (symbols). This means that before using a continuous feature vector, we must convert it to discrete values. A common way to do that is through clustering. The problem is that for the most of real-life situations we do not know the best number of clusters, what makes it necessary to explore different numbers of clusters using traditional clustering methods such as the K-means algorithm [43] and its variants. In this light, clustering can become a trial-and-error work. Besides, its result may not be very promising especially when the number of clusters is large and not easy to estimate.

Unsupervised feature selection emerges as a clever solution to this problem. The literature contains several studies on feature selection for supervised learning, but only recently, the feature selection for unsupervised learning has been investigated [44, 45]. The objective in unsupervised feature selection is to search for a subset of features that best

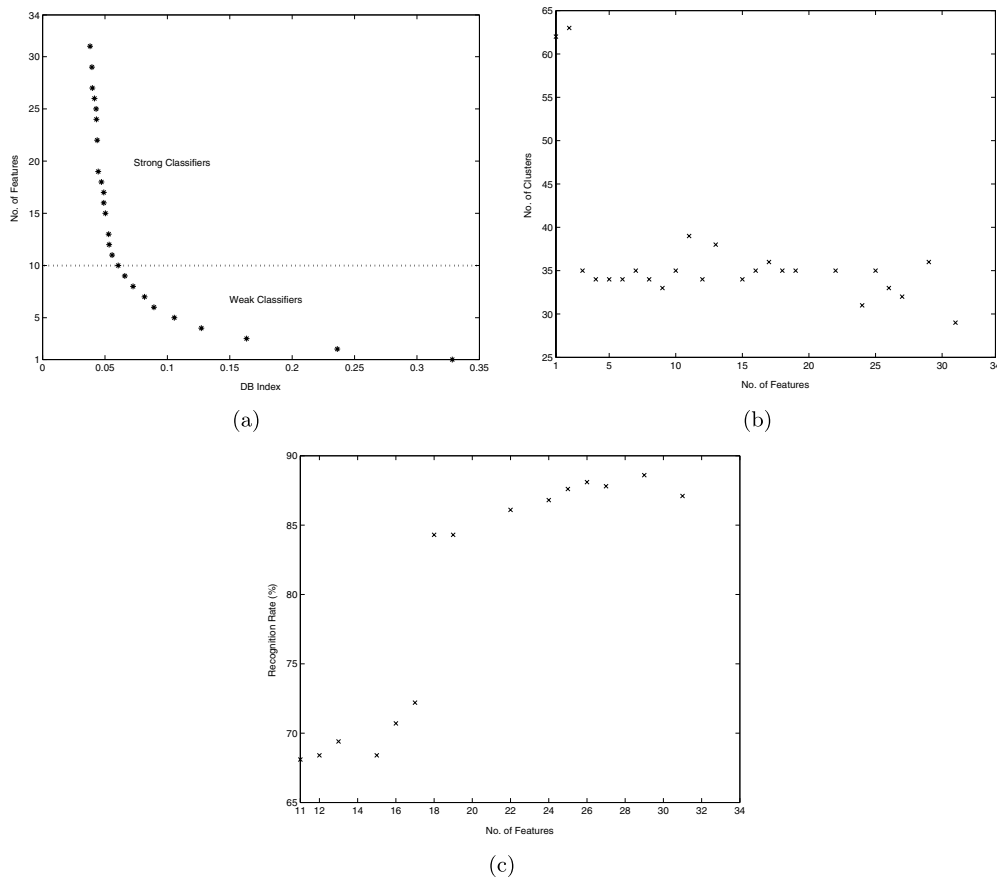


Fig. 9 **a** Pareto-optimal front found during feature selection divided into two classes of classifiers, **b** Relationship between the number of clusters and the number of features and **c** relationship between the recognition rate and the number of features

uncovers “natural” groupings (clusters) from data according to some criterion. In this way, we can avoid the manual process of clustering and find the most discriminative features in the same time. Hence, we will have at the end a more compact and robust high-level representation (symbols).

In the above context, unsupervised feature selection also presents a multi-criterion optimization function, where the objective is to find compact and well separated hyper-spherical clusters in the feature subspaces. Differently of the supervised feature selection, here the criteria optimized by the algorithm are a validity index and the number of features. [46].

In order to measure the quality of clusters during the clustering process, we have used the Davies-Bouldin (DB)-index [47] over 80,000 feature vectors extracted from the training set of 9,500 words. To make such an index suitable for our problem, it must be normalized by the number of selected features. This is due to the fact that it is based on geometric distance metrics and therefore, it is not directly applicable here because it is biased by the dimensionality of the space, which is variable in feature selection problems.

We have noticed that the value of DB index decreases as the number of features increases. We have correlated this effect with the normalization of DB-index by the number of features. In order to compensate this, we have considered

as second objective the minimization of the number of features. In this case, one feature must be set at least. Figure 9 depicts the Pareto-optimal front found after the search, the relationship between the number of clusters and number of features and the relationship between the recognition rate on the validation set and the number of features.

Like in the supervised context, here we also divided the classifiers of the Pareto into classes. In this case, we have realized that those classifiers with very few features are not selected to compose the ensemble, and therefore, just the classifiers with more than 10 features were used into the second level of search. In Sect. 7.2 we discuss this issue in more detail. Figure 9 shows the Pareto-optimal front where a line divides the classifiers into two different groups: weak (less than 10 features) and strong (more than 10 features). The way of choosing the best ensemble is exactly the same as introduced in Sect. 6.1.2.

7 Experimental results

All experiments in this work were based on a single-population master-slave MOGA. In this strategy, one master node executes the genetic operators (selection, crossover and mutation), and the evaluation of fitness is distributed

among several slave processors. We have used a Beowulf cluster with 17 (one master and 16 slaves) PCs (1.1 Ghz CPU, 512 Mb RAM) to execute our experiments.

The following parameter settings were employed in both levels: population size = 128, number of generations = 1,000, probability of crossover = 0.8, probability of mutation = $1/L$ (where L is the length of the chromosome), and niche distance (σ_{share}) = [0.25, 0.45]. The length of the chromosome in the first level is the number of components in the feature set (see Table 1), while in the second level is the number of classifiers picked from the Pareto-optimal front in the previous level.

In order to define the probabilities of crossover and mutation, we have used the one-max problem, which is probably the most frequently-used test function in research on genetic algorithms because of its simplicity [48]. This function measures the fitness of an individual as the number of bits set to one on the chromosome. We have used a standard genetic algorithm with a single-point crossover and the maximum generations of 1,000. The fixed crossover and mutation rates are used in a run, and the combination of the crossover rates 0.0, 0.4, 0.6, 0.8 and 1.0 and the mutation rates of $0.1/L$, $1/L$ and $10/L$, where L is the length of the chromosome. The best results were achieved with $P_c = 0.8$ and $P_m = 1/L$. Such results confirmed the values reported by Miki et al in [49]. The parameter σ_{share} was determined initially through Eq. (2) and then tuned empirically.

We have run each experiment at least 10 times and we have noticed that the best solution was stable over the runs. Using the aforementioned cluster the experiments in the supervised context took about 2h, while in the unsupervised took about 8h. The latter is more time consuming because of the clustering algorithm used during feature selection.

7.1 Experiments in the supervised context

Once all parameters have been defined, the first step, as described in Section 6.1.1, consists of performing feature selection for a given feature set. As depicted in Fig. 7, this procedure produces quite a large number of classifiers, which should be trained for use in the second level. After some experiments, we found out that the second level always chooses “strong” classifiers to compose the ensemble. Thus, in order to speed up the training process and the second level of search as well, we decide to train and use in the second level just the “strong” classifiers. This decision was made after we realized that in our experiments the “weak” and “medium” classifiers did not cooperate with the ensemble at all. To train such classifiers, the same databases reported in Sect. 5.1 were used. Table 3 summarizes the “strong” classifiers produced by the first level for the three feature sets we have considered.

Considering for example the feature set CC_{sc} , the first level of the algorithm provided 81 “strong” classifiers which have the number of features ranging from 24 to 125 and

Table 3 Summary of the classifiers produced by the first level

Feature set	Number of classifiers	Range of features	Range of rec. rates (%)
CC_{sc}	81	24–125	90.5–99.1
DDD_{sc}	54	30–84	90.6–98.1
EM_{sc}	78	35–113	90.5–97.0

recognition rates ranging from 90.5% to 99.1% on $TSDB_{sc}$. This shows the great diversity of the classifiers produced by the feature selection method. Based on the classifiers reported in Table 3 we define four sets of base classifiers as follows: $S_1 = \{CC_{sc0}, \dots, CC_{sc80}\}$, $S_2 = \{DDD_{sc0}, \dots, DDD_{sc53}\}$, $S_3 = \{EM_{sc0}, \dots, EM_{sc77}\}$, and $S_4 = \{S_1 \cup S_2 \cup S_3\}$. All these sets could be seen as ensembles, but in this work we reserve the word ensemble to characterize the results yielded by the second-level of the algorithm. In order to assess the objective functions of the second-level of the algorithm (generalization of the ensemble and diversity) we have used the validation set ($VLDB1_{sc}$).

Like the first level, the second one also generates a set of possible solutions which are the trade-offs between the generalization of the ensemble and its diversity. Thus the problem now lies in choosing the most accurate ensemble among all. Figure 10 depicts the variety of ensembles yielded by the second-level of the algorithm for the four sets of base classifiers being considered. The number over each point stands for the number of classifiers in the ensemble. In order to decide which ensemble to choose we validate the Pareto-optimal front using $VLDB2_{sc}$, which was not used so far. Since we are aiming at performance, the direct choice will be the ensemble that provides better generalization on $VLDB2_{sc}$. Table 4 summarizes the best ensembles produced for the four sets of base classifiers and their performance at zero-rejection level on the test set. For facility, we reproduce in this table the results of the original classifiers.

We can notice from Table 4 that the ensembles and base classifiers have very similar performance at zero-rejection level. On the other hand, Fig. 11 shows that the ensembles respond better for error rates fixed at very low levels than single classifiers. The most expressive result was achieved for the ensemble S_3 , which attains a reasonable performance at zero-rejection level but performs very poorly at low error rates. In such a case, the ensemble of classifiers brought an improvement of about 8%. We have noticed that the ensemble reduces the high outputs of some outliers so that

Table 4 Performance of the ensembles on the test set

Feature set	Number of classifiers	Rec. rate (%)	
		Zero-rejection level	Original classifiers
S_1	4	99.22	99.13
S_2	4	98.18	98.17
S_3	7	97.10	97.04
S_4	24	99.25	

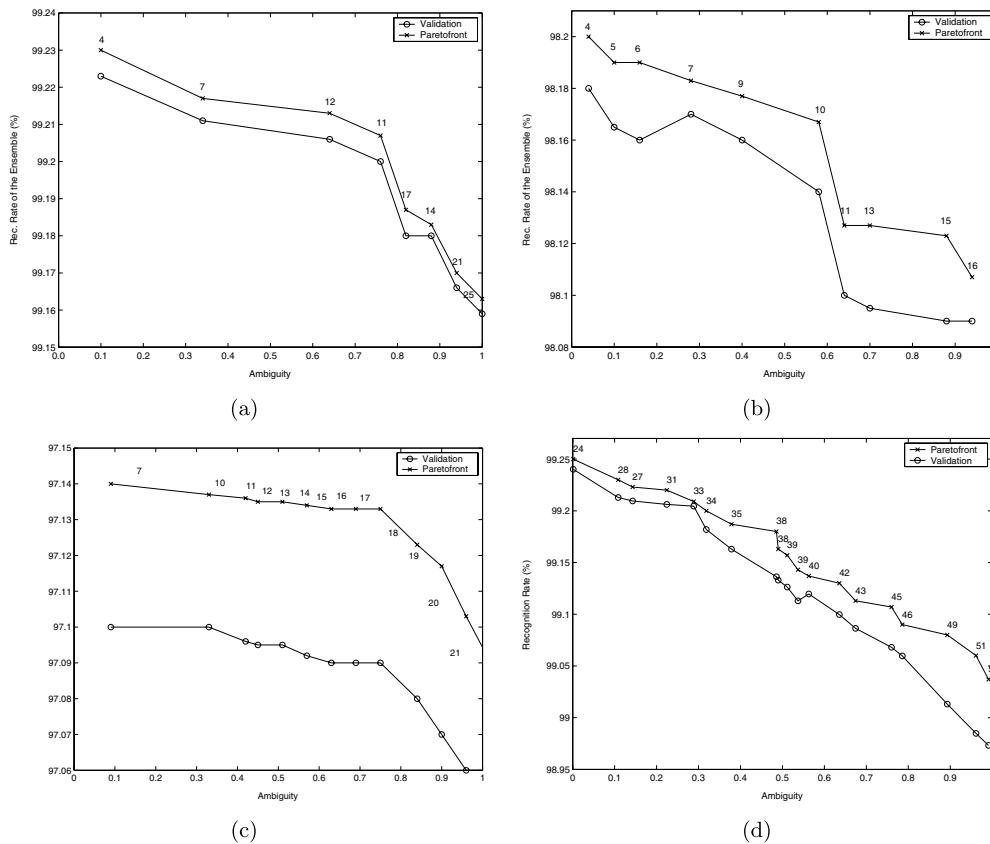


Fig. 10 The Pareto-optimal front produced by the second-level MOGA: S_1 **a**, S_2 **b**, S_3 **c**, and S_4 **d**

the threshold used for rejection can be reduced and consequently the number of samples rejected is reduced. Thus, aiming for a small error rate we have to consider the important role of the ensemble.

Regarding the ensemble S_4 , we can notice that it achieves a performance similar to S_1 at zero-rejection level (see Table 4). Besides, it is composed of 24 classifiers, against four of S_1 . The fact worth noting though, is the performance of S_4 at low error rates. For the error rate fixed at 1% it reached 95.0% against 93.5% of S_1 . S_4 is composed of 14, 6, and 4 classifiers from S_1 , S_2 , and S_3 , respectively. At this point, we would like to stress that the ensembles can improve results even when just one feature set is available, as depicted in Fig. 11a, b, and c. However, when more original classifiers are available, it is expected that the algorithm builds better ensembles. This is what we observe in Fig. 11d for the ensemble S_4 .

To conclude the experiments in the supervised context, we present in Fig. 12 some examples of misclassification generated by the ensembles, where those signed with a square would represent mislabelling. We can observe that most of this errors are pretty difficult to recognize, with the exception of digits “1”, which are not that difficult, but barely appear in the training set of NIST.³

³ This style of handwriting is very similar to that found in the Brazilian database.

7.2 Experiments in the unsupervised context

The experiments in the unsupervised context follow the same vein of the supervised one. As discussed in Sect. 6.2.1, the main difference lies in the way the feature selection is carried out. In spite of that, we can observe that the number of classifiers produced during unsupervised feature selection is quite large as well. In light of this, we have applied the same strategy of dividing the classifiers into groups (see Fig. 9). After some experiments, we found out that the second level always chooses “strong” classifiers to compose the ensemble. Thus, in order to speed up the training process and the second level of search as well, we decide to train and use in the second level just “strong” classifiers. To train such classifiers, the same databases reported in Sect. 5.2 were considered. Table 5 summarizes the “strong” classifiers (after training) produced by the first level for the three feature sets we have considered.

Considering for example the feature set $CCuc$, the first level of the algorithm provided 15 “strong” classifiers which have the number of features ranging from 10 to 32 and recognition rates ranging from 68.1 to 88.6% on $VLDB1uc$. This shows the great diversity of the classifiers produced by the feature selection method. Based on the classifiers reported in Table 5 we define four sets of base classifiers as follows:

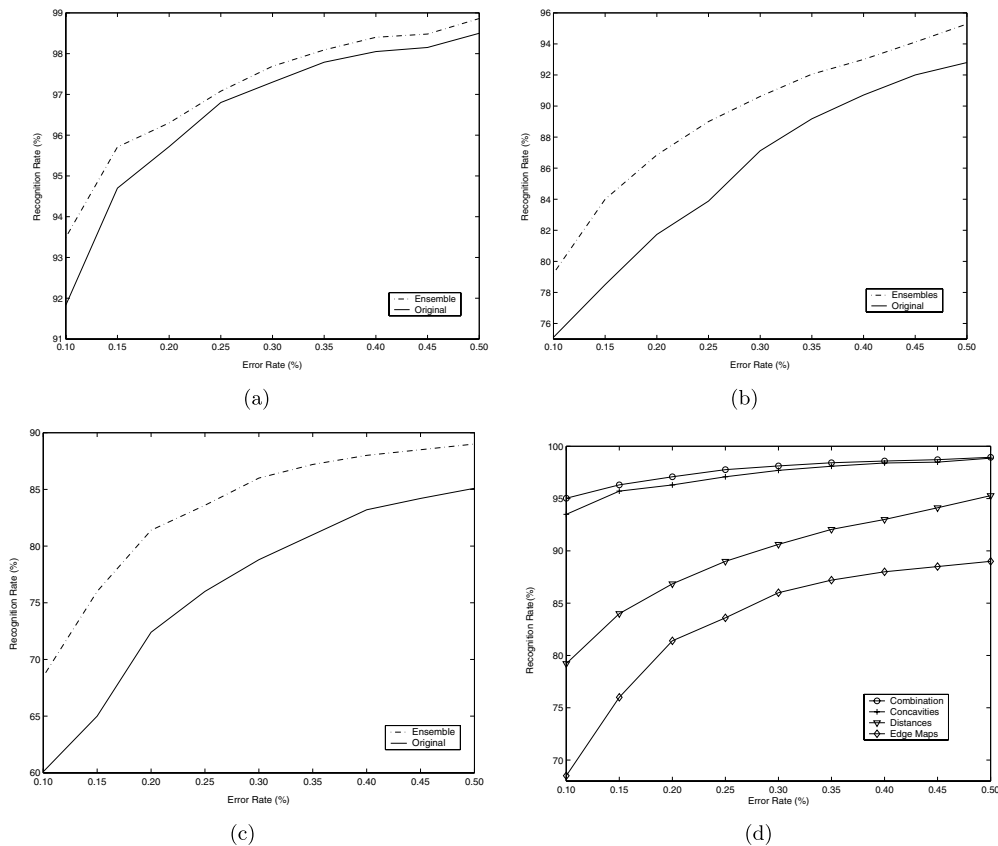


Fig. 11 Improvements yielded by the ensembles: S_1 a, S_2 b, S_3 c, and S_4 d

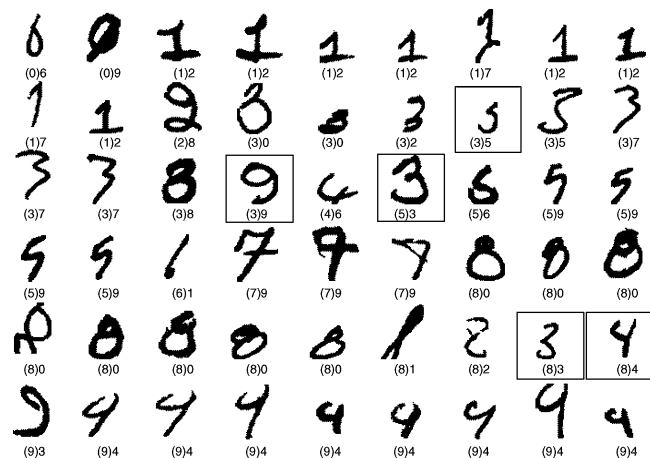


Fig. 12 Examples of misclassification (the correct label is the one in parentheses)

$F_1 = \{CCuc_0, \dots, CCuc_{14}\}$, $F_2 = \{DDD32uc_0, \dots, DD D32uc_{20}\}$, $F_3 = \{DDD64uc_0, \dots, DDD64uc_{49}\}$, and $F_4 = \{F_1 \cup F_2 \cup F_3\}$.

Figure 13 depicts the variety of ensembles yielded by the second-level of the algorithm for the four sets of base classifiers. The number over each point stands for the number of classifiers in the ensemble. Like in the previous experiments,

Table 5 Summary of the classifiers produced by the first level

Feature set	Number of classifiers	Range of features	Range of codebook	Range of rec. rates (%)
<i>CCuc</i>	15	10–32	29–39	68.1–88.6
<i>DDD32uc</i>	21	10–31	20–30	71.7–78.0
<i>DDD64uc</i>	50	10–64	52–80	60.6–78.2

Table 6 Comparison between ensembles and original classifiers

Base classifiers	Number of Rec. rate	(%) classifiers	Original feature set	Rec. rate (%)
F_1	10	89.2	CC	86.1
F_2	15	80.2	DDD32	73.0
F_3	36	80.7	DDD64	64.5
F_4	45	90.2		

the second validation set (*VLDB2uc*) was used to select the best ensemble. After selecting the best ensemble the final step is to assess them on the test set. Table 6 summarizes the performance of the ensembles on the test set. For the sake of comparison, we reproduce in Table 6 the results presented in Table 2.

Figure 13 also shows the performance of the ensembles generated with all base classifiers available, i.e., Ensemble F_4 . Like in the previous experiments (supervised context), the result achieved by the ensemble F_4 shows the ability of

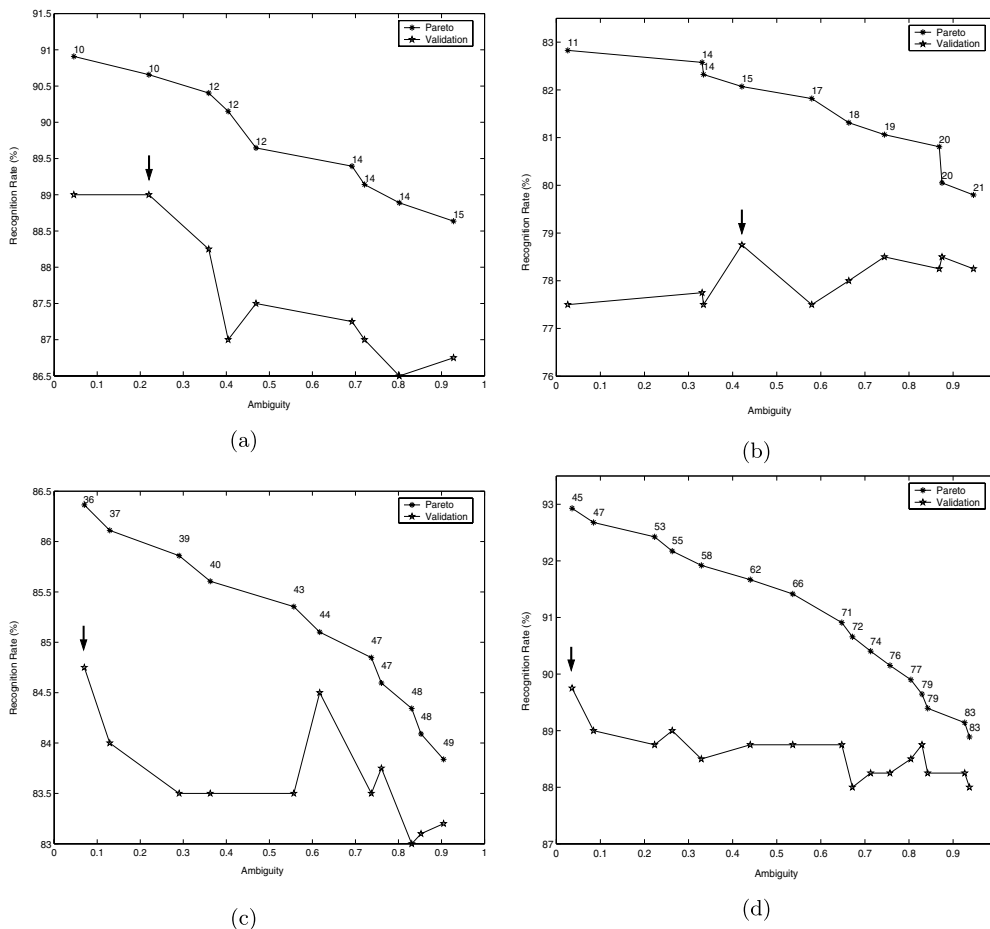


Fig. 13 The Pareto-optimal front (and validation curves where the best solutions are highlighted with an arrow) produced by the second-level MOGA: F_1 **a**, F_2 **b**, F_3 **c**, and F_4 **d**

the algorithm in finding good ensembles when more base classifiers are considered. The ensemble F_4 is composed of 9, 11, and 25 classifiers from F_1 , F_2 , and F_3 , respectively.

In light of this, we decided to introduce a new feature set, which, based on our experience, has a good discrimination power when combined with other features such as concavities. This feature set, which we call “global features”, is composed of primitives such as ascenders, descenders, and loops. The combination of these primitives plus a primitive that determines whether a grapheme does not contain ascender, descender, and loop produces a 20-symbol alphabet. For more details, see [50]. In order to train the classifier with this feature set, we have used the same databases described in Sect. 5.2. The recognition rates at zero-rejection level are 86.1% and 87.2% on validation and testing sets, respectively. This performance compares with the *CCuc* classifier.

Since we have a new base classifier, our sets of base classifiers must be modified to cope with it. Thus, $F_{1G} = \{F_1 \cup G\}$, $F_{2G} = \{F_2 \cup G\}$, $F_{3G} = \{F_3 \cup G\}$, and $F_{4G} = \{F_1 \cup F_2 \cup F_3 \cup G\}$. In such cases, G stands for the classifier trained with global features. Table 7 summarizes the ensembles found using these new sets of base classifiers.

It is worthy of remark the reduction of the size of the teams. This shows the ability of the algorithm in finding not just diverse but also uncorrelated classifiers to compose the ensemble [51]. Besides, it corroborates our claim that the classifier G when combined with other features bring an improvement to the performance.

In Fig. 14 we compare the error-reject trade-offs of the ensembles reported in Table 7. Like the results at zero-rejection level, the improvement observed here also are quite impressive. Table 7 shows that F_{1G} and F_{4G} reach similar results on the test set at zero-rejection level, however, F_{1G} contains just two classifiers against 23 of F_{4G} . On the other

Table 7 Performance of the ensembles with global features

Base classifiers	Number of classifiers	Rec. rate (%) testing
F_{1G}	2	92.2
F_{2G}	2	89.7
F_{3G}	7	85.5
F_{4G}	23	92.0

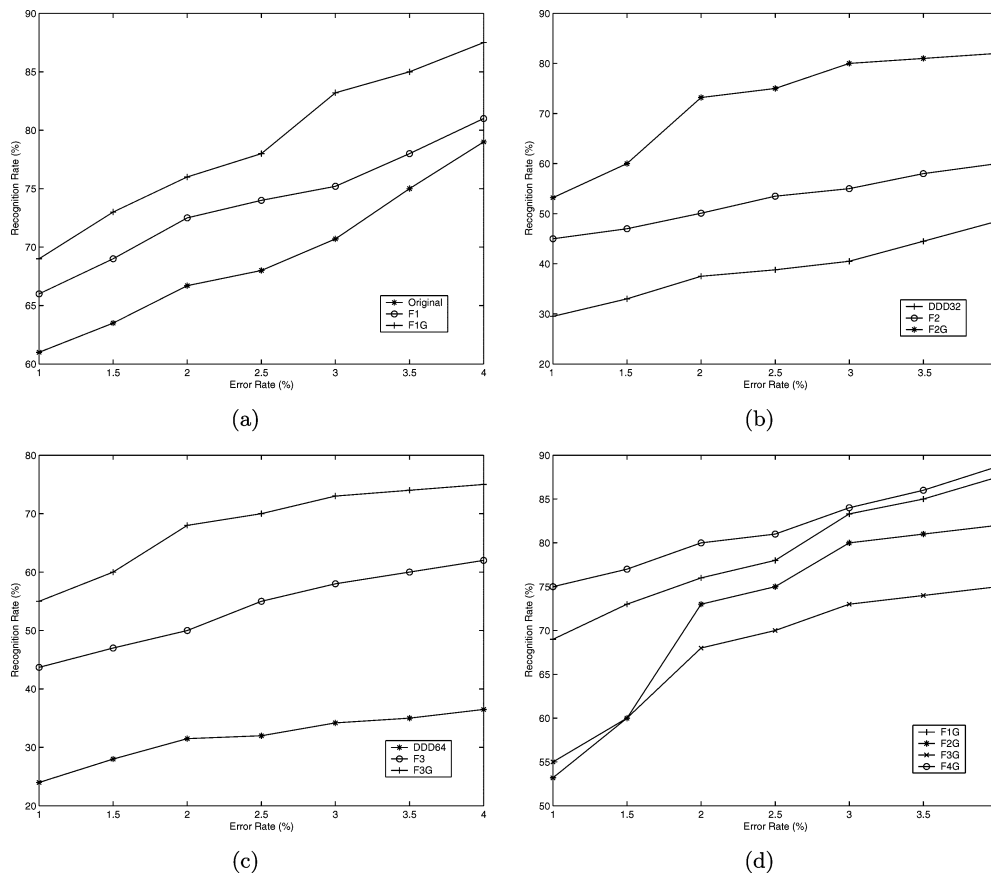


Fig. 14 Improvements yielded by the ensembles: F_1 a, F_2 b, F_3 c, and comparison among all ensembles d

hand, the latter features a slightly better error-reject trade-off in the long run (Fig. 14d).

Based on the experiments reported so far we can affirm that the unsupervised feature selection is a good strategy to generate diverse classifiers. This is made very clear in the experiments regarding the feature set DDD64. In such a case, the original classifier has a poor performance (about 65% on the test set), but when it is used to generate the set of base classifiers, the second-level MOGA was able to produce a good ensemble by maximizing the performance and the ambiguity measure. Such an ensemble of classifiers brought an improvement of about 15% in the recognition rate at zero-rejection level. Figure 15 some examples of well-classified and misclassified images of our database.




Fig. 15 Examples of a well-classified images and b misclassified images (the correct label is the one in parentheses)

8 Discussion

The results obtained here attest that the proposed strategy is able to generate a set of good classifiers in both supervised and unsupervised contexts. To better evaluate our results, we have used two traditional ensemble methods (Bagging and Boosting) in the supervised context. Figure 16 reports the results. As we can see, the proposed methodology achieved better results, especially when considering very low error rates. It also can be observed from this figure that the gap among ensemble feature selection and bagging/boosting gets bigger as the discriminative power of the feature sets gets lower. However to verify whether this holds, experimentation with other feature sets would be necessary.

Diversity is an issue that deserves some attention when discussing ensemble of classifiers. As we have mentioned before, some authors advocated that diversity does not help at all. In our experiments, most of the time, the best ensembles of the Pareto-optimal also were the best for the unseen data. This could lead one to agree that diversity is not important when building ensembles, since even using a validation set the selected team is always the most accurate and with less diversity.

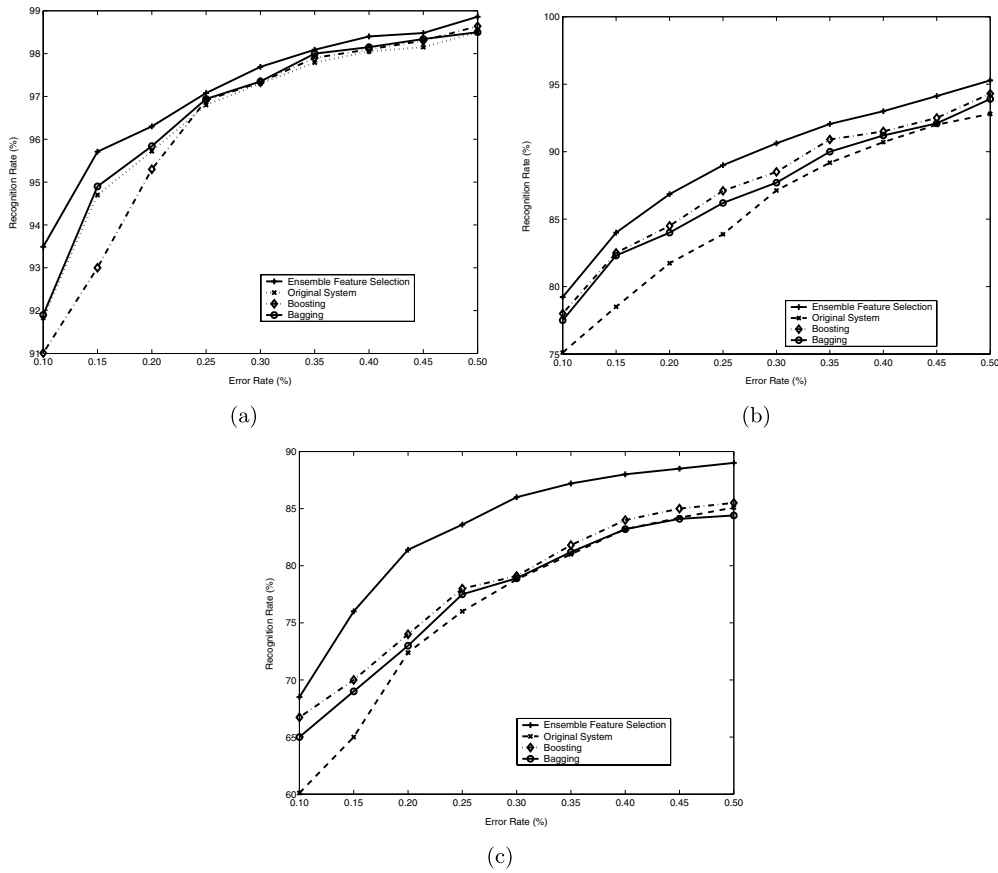


Fig. 16 Comparison among ensemble feature selection, bagging, and boosting for the tree feature set used in the supervised context: **a** CC_{sc} , **b** DDD_{sc} , and **c** EM_{sc}




Fig. 17 Benefits of using diversity. **a** Population (classifiers) of the last generation of the GA. **b** Classifiers found by the MOGA

However, if we look carefully the results, we will observe that there are cases where the validation curve does not have the same shape of the Pareto-optimal. In such cases diversity is very useful to avoid selecting overfitted solutions.

One can argue that using a single GA and considering the entire final population, perhaps the similar solutions found

in the Pareto-optimal produced by the MOGA will be there. To show that it does not happen, we have carried out some experiments with a single GA where the fitness function was the maximization of the ensembles accuracy. Since a single-objective optimization algorithm searches for an optimum solution, it is natural to expect that it will converge towards the fittest solution, hence, the diversity of solutions presented in

the Pareto-optimal is not present in the final population of the single genetic algorithm.

To illustrate that, we present the results we got using a GA to find ensemble in F_2 (unsupervised context). The parameters used here are the same we have used for the MOGA (Sect. 7). Figure 17a plots all the classifiers found in the last population of the genetic algorithm. For the sake of comparison we reproduce Fig. 13b in Fig. 17b. As we can see, the population is very homogeneous and it converged, as expected, towards the most accurate ensemble.

Some attempts in this direction were made by Optiz [4]. He combined accuracy and diversity through the weighted-sum approach. As stated somewhere, when dealing with this kind of combination, one should deal with problems such as scaling and sensitivity towards the weights. We believe that our strategy offers a clever way to find the ensemble using genetic algorithms.

9 Conclusion

We have described a methodology for ensemble creation underpinned on the paradigm “overproduce and choose”. It takes two levels of search where the first level overproduces a set of classifiers by performing feature selection while the second one chooses the best team of classifiers.

The feasibility of the strategy was demonstrated through comprehensive experiments carried out in the context of handwriting recognition. The idea of generating classifiers through feature selection was proved to be successful in both supervised and unsupervised contexts. The results attained in both situations and using different feature sets and base classifiers demonstrated the efficiency of the proposed strategy by finding powerful ensembles, which succeed in improving the recognition rates for classifiers working with a very low error rates. Such results compare favorably to traditional ensemble methods such as Bagging and Boosting.

Finally we have addressed the issue of using diversity to build ensembles. As we have seen, using diversity jointly with the accuracy of the ensemble as selection criterion might be very helpful to avoid choosing overfitted solutions. Our results certainly brings some contribution to the field, but this still is an open problem.

References

- Hansen, L., Salomon, O.: Neural network ensembles. *IEEE Trans. Pattern Anal. Mach. Intell.* **12**(10), 993–1001 (1990)
- Hashem, S.: Optimal linear combinations of neural networks. *Neural Networks* **10**(4), 599–614 (1997)
- Krogh, A., Vedelsby, J.: Neural networks ensembles, cross validation, and active learning. In: Tesauro, G. et al. (eds.). *Advances in Neural Information Processing Systems 7*, pp. 231–238. MIT Press (1995)
- Optiz, D.W.: Feature selection for ensembles. In: *Proceedings of the 16th International Conference on Artificial Intelligence*, pp. 379–384 (1999)
- Ho, T.K.: The random subspace method for constructing decision forests. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**(8), 832–844 (1998)
- Last, M., Bunke, H., Kandel, A.: A feature-based serial approach to classifier combination. *Pattern Anal. Appl.* **5**, 385–398 (2002)
- Tsybmal, A., Puuronen, S., Patterson, D.W.: Ensemble feature selection with the simple Bayesian classification. *Inf. Fusion* **4**, 87–100 (2003)
- Gerra-Salcedo, C., Whitley, D.: Genetic approach to feature selection for ensemble creatin. In: *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 236–243. Orlando-USA (1999)
- De Stefano, C., Cioppa, A.D., Marcelli, A.: Exploiting reliability for dynamic selection of classifiers by means of genetic algorithms. In: *Proceedings of the 7th International Conference on Document Analysis and Recognition*, pp. 671–675. Edinburgh-Scotland IEEE Computer Society (2003)
- Emmanouilidis, C., Hunter, A., MacIntyre, J.: A multiobjective evolutionary setting for feature selection and a commonality-based crossover operator. In: *Proceedings of the Congress on Evolutionary Computation*, vol. 1, pp. 309–316 (2000)
- Oliveira, L.S., Sabourin, R., Bortolozzi, F., Suen, C.Y.: A methodology for feature selection using multi-objective genetic algorithms for handwritten digit string recognition. *Int. J. Pattern Recog. Artif. Intell.* **17**(6), 903–930 (2003)
- Siedlecki, W., Sklansky, J.: A note on genetic algorithms for large scale on feature selection. *Pattern Recog. Lett.* **10**, 335–347 (1989)
- Kudo, M., Sklansky, J.: Comparison of algorithms that select features for pattern classifiers. *Pattern Recog.* **33**(1), 25–41 (2000)
- Partridge, D., Yates, W.B.: Engineering multiversion neural-net systems. *Neural Comput.* **8**(4), 869–893 (1996)
- Giacinto, G., Roli, F.: Design of effective neural network ensemble for image classification purposes. *Image Vision Comput. J.* **9–10**, 697–705 (2001)
- Kuncheva, L.I., Whitaker, C.J.: Measures of diversity in classifier ensembles. *Mach. Learn.* **51**, 181–207 (2003)
- Ruta, D., Gabrys, B.: Analysis of the correlation between majority voting error and the diversity measures in multiple classifier systems. In: *Proceedings of the 4th International Symposium on Soft Computing*. Paisley, UK (2001)
- Ruta, D.: Multilayer selection-fusion model for pattern classification. In: *Proceedings of the IASTED Artificial Intelligence and Application Conference*. Innsbruck, Austria (2004)
- Kuncheva, L.: That elusive diversity in classifier ensembles. In: *Proceedings of the ibPRIA, LNCS 2652*, pp. 1126–1138. Mallorca, Spain (2003)
- Jain, A.K., Duin, R.P.W., Mao, J.: Statistical pattern recognition: A review. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(1), 4–37 (2000)
- Kuncheva, L., Bezdek, J.C., Duin, R.P.W.: Decision templates for multiple classifier fusion: An experimental comparison. *Pattern Recog.* **34**(2), 299–314 (2001)
- Efron, B., Tibshirani, R.: *An Introduction to the Bootstrap*. Chapman and Hall (1993)
- Breiman, L.: Stacked regressions. *Mach. Learn.* **24**(1), 49–64 (1996)
- Freund, Y., Schapire, R.: Experiments with a new boosting algorithm. In: *Proceedings of the 13th International Conference on Machine Learning*, pp. 148–156. Bary-Italy (1996)
- Tumer, K., Oza, N.C.: Input decimated ensembles. *Pattern Anal. Appl.* **6**, 65–77 (2003)
- Kuncheva, L., Jain, L.C.: Designing classifier fusion systems by genetic algorithms. *IEEE Trans. Evol. Comput.* **4**(4), 327–336 (2000)
- Gunter, S., Bunke, H.: Creation of classifier ensembles for handwritten word recognition using feature selection algorithms. In: *Proceedings of the 8th IWFHR*, pp. 183–188. Niagara-on-the-Lake, Canada (2002)

28. Deb, K.: *Multi-Objective Optimization using Evolutionary Algorithms*. Wiley, 2nd edn. April (2002)
29. Srinivas, N., Deb, K.: Multiobjective optimization using nondominated sorting in genetic algorithms. *Evol. Comput.* **2**(3), 221–248 (1995)
30. Deb, K., Goldberg, D.E.: An investigation of niche and species formation in genetic function. In: *Proceedings of the 3rd International Conference on Genetic Algorithms*, pp. 42–50 (1989)
31. Oliveira, L.S., Sabourin, R., Bortolozzi, F., Suen, C.Y.: Automatic recognition of handwritten numerical strings: A recognition and verification strategy. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**(11), 1438–1454 (2002)
32. Oh, I.-S., Suen, C.Y.: Distance features for neural network-based recognition of handwritten characters. *Int. J. Doc. Anal. Recog.* **1**(2), 73–88 (1998)
33. Chim, Y.C., Kassim, A.A., Ibrahim, Y.: Dual classifier system for handprinted alphanumeric character recognition. *Pattern Anal. Appl.* **1**(3), 155–162 (1998)
34. Bishop, C.M.: *Neural Networks for Pattern Recognition*. Oxford University Press, Oxford, U.K. (1995)
35. Fumera, G., Roli, F., Giacinto, G.: Reject option with multiple thresholds. *Pattern Recog.* **33**(12), 2099–2101 (2000)
36. Oliveira, J.J. Jr., Carvalho, J.M., Freitas, C.O.A., Sabourin, R.: Evaluating NN and HMM classifiers for handwritten word recognition. In: *Proceedings of the 15th Brazilian Symposium on Computer Graphics and Image Processing*, pp. 210–217. IEEE Computer Society Fortaleza, Brazil (2002)
37. Morita, M., El Yacoubi, A., Sabourin, R., Bortolozzi, F., Suen, C.Y.: Handwritten month word recognition on Brazilian bank cheques. In: *Proceedings of the 6th ICDAR*, pp. 972–976 (2001)
38. John, G., Kohavi, R., Pfleger, K.: Irrelevant features and the subset selection problems. In: *Proceedings of the 11th International Conference on Machine Learning*, pp. 121–129 (1994)
39. Moody, J., Utans, J.: Principled architecture selection for neural networks: Application to corporate bond rating prediction. In: Moody, J., Hanson, S.J., Lippmann, R.P. (eds.) *Advances in Neural Information Processing Systems 4*. Morgan Kaufmann (1991)
40. Yuan, H., Tseng, S.S., Gangshan, W., Fuyan, Z.: A two-phase feature selection method using both filter and wrapper. In: *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*, vol. 2, pp. 132–136 (1999)
41. Kuncheva, L.I., Whitaker, C.J.: Ten measures of diversity in classifier ensembles: limits for two classifiers. In: *Proceedings of the IEE Workshop on Intelligent Sensor Processing*, pp. 1–10 (2001)
42. Kittler, J., Hatef, M., Duin, R., Matas, J.: On combining classifiers. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**(3), 226–239 (1998)
43. Huang, X.D., Ariki, Y., Jack, M.A.: *Hidden Markov Models for Speech Recognition*. Edinburgh University Press (1990)
44. Dy, J.G., Brodley, C.E.: Feature subset selection and order identification for unsupervised learning. In: *Proceedings of the 17th International Conference on Machine Learning*. Stanford University, CA, July (2000)
45. Kim, Y.S., Street, W.N., Menczer, F.: Feature selection in unsupervised learning via evolutionary search. In: *Proceedings of the 6th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 365–369 (2000)
46. Morita, M., Sabourin, R., Bortolozzi, F., Suen, C.Y.: Unsupervised feature selection using multi-objective genetic algorithms for handwritten word recognition. In: *Proceedings of the 7th International Conference on Document Analysis and Recognition*, pp. 666–670. IEEE Computer Society, Edinburgh-Scotland (2003)
47. Davis, L.: *Handbook on Genetic Algorithms*. Van Nostrand Reinhold (1991)
48. Cantu-Paz, E.: *Efficient and Accurate Parallel Genetic Algorithms*. Kluwer Academic Publishers (2000)
49. Miki, M., Hiroyasu, T., Kaneko, K., Hatanaka, K.: A parallel genetic algorithm with distributed environment scheme. In: *Proceedings of the International Conference on System, Man, and Cybernetics*, vol. 1, pp. 695–700 (1999)
50. Morita, M., Oliveira, L.S., Sabourin, R., Bortolozzi, F., Suen, C.Y.: An HMM-MLP hybrid system to recognize handwritten dates. In: *Proceedings of the International Joint Conference on Neural Networks*, pp. 1–6. IEEE Computer Society, Honolulu, USA (2002)
51. Tumer, K., Ghosh, J.: Error correlation and error reduction in ensemble classifiers. *Connect. Sci.* **8**(3–4), 385–404 (1996)