

ADVANCES IN FEATURE SELECTION METHODOLOGY: AN OVERVIEW OF RECENT UTIA RESULTS

PETR SOMOL, JANA NOVOVIČOVÁ AND PAVEL PUDIL

The paper gives an overview of feature selection techniques in statistical pattern recognition with particular emphasis to methods developed by our UTIA research team throughout recent years. Besides discussing the advances in methodology it attempts to put them into a taxonomical framework. The methods discussed include the latest variants of the optimal algorithms, enhanced sub-optimal techniques and the simultaneous semi-parametric probability density function modelling and feature space selection method. Some related issues are illustrated on real data by means of the Feature Selection Toolbox software.

Keywords: feature selection, branch & bound, sequential search, mixture model

AMS Subject Classification: 62H30, 62G05, 68T10

1. INTRODUCTION

The application of pattern recognition and machine learning in the real-world domain often encounters data problems caused by the high dimensionality of the input space. Dimensionality reduction refers to the task of finding low dimensional representations for high-dimensional data. More formally, given a set of n real D -dimensional vectors $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, $\mathbf{x}_i \in \mathcal{R}^D$, we define dimensionality reduction as the task of finding a set of corresponding low-dimensional representatives $\{\mathbf{y}_1, \dots, \mathbf{y}_n\}$, $\mathbf{y}_i \in \mathcal{R}^d$, that preserve the similarity structure in the input data as well as possible. Dimensionality reduction (DR) is an important step in data preprocessing in pattern recognition and machine learning applications. It is sometimes the case that such tasks as classification or approximation of the data represented by so called feature vectors, can be carried out in the reduced space more accurately than in the original space.

There are two main ways of doing DR depending on the resulting features: DR by *feature selection* (FS) and DR by *feature extraction* (FE). FS approach does not attempt to generate new features, but to try to select the "best" ones from the original set of features. FE approach defines a new feature vector space in which each new feature is obtained by combinations or transformations of the original features. FS leads to savings in measurements cost since some of the features are discarded and the selected features retain their original physical interpretation. In addition, the retained features may be important for understanding the physical process that generates the feature vectors. On the other hand, transformed features generated by

feature extraction may provide a better discriminative ability than the best subset of given features, but these new features may not have a clear physical meaning.

We are honored to acknowledge that the pioneer research in this direction was carried out already as far back as in 1967 by Albert Perez and his co-workers. In the framework of a broader research project "Decision processes" (in medical care) they started to investigate problems of DR. At that time they used the term "reduction of symptom complex". Perez has shown that an intuitive reduction can lead to substantial loss of information and proposed to use mathematical methods and computers to optimize the solution. The theoretical background was formed mainly by Perez's own theoretical results like α -information values. As a result he and his co-workers developed several variants of algorithm TIBIS (Theory of Information in Biological Systems). The algorithms used either the forward (including) strategy or the backward (eliminating) strategy. The first series of PC codes was implemented in years 1967-71 on MINSK 22 mainframe, in 1972-73 a more advanced and sophisticated version was implemented on IBM 370 in Fortran. Obviously, the implementation was restricted by hardware limitation of those early years, however, for its time it was a remarkable achievement. All the algorithms yielded subsets of symptoms from original symptom complex - thus in the present terminology carried out feature selection.

This paper is limited to the feature selection approach to dimensionality reduction. We introduce concepts and algorithms of feature selection, present a collection of feature selection algorithms for decision problems of classification type developed by our research team, compare different algorithms and provide guidelines in selecting the feature selection algorithms. Many of recent feature selection methods have been implemented in our Feature Selection Toolbox software. Some trends in feature selection research and development are presented.

2. FEATURE SUBSET SELECTION

Let X_D be the set that comprises all the D original features, and \mathcal{X}_d the set of all possible subsets of size d , where d represents the desired number of features. Let X be a subset of X_D , with d features. $J(X)$ is a function that determines how good the subset X is, by a certain criterion. Without any loss of generality, let us assume that a higher value of J indicates a better feature subset.

The problem of feature subset selection is defined by searching the subset $\tilde{X}_d \subset X_D$ composed by d features, that satisfies:

$$J(\tilde{X}_d) = \max_{X \in \mathcal{X}_d} J(X). \quad (1)$$

An optimal subset is always relative to a certain evaluation criterion (i.e., an optimal subset chosen using one evaluation function may not be the same as that using another evaluation function).

A feature selection process consists of four basic steps: *subset generation*, *subset evaluation*, *stopping criterion*, and *result validation* [13]. Selection strategies and evaluation criteria are two dominating factors in designing a feature selection algorithm.

Liu and Yu [13] provide a comprehensive overview of various aspects of feature selection.

2.1. Categorization of FS methods

Different methods have been developed and used for feature subset selection in pattern recognition and machine learning, using different search strategies and evaluation functions. According to Liu and Yu [13], the basic approach to subset generation is to build up a subset of required number of features starting with the empty set and successively add features (*forward* search), or to start with a full set and remove features successively (*backward* search), or to start with both ends and add and remove features simultaneously (*bi-directional* search). Search may also start with a randomly selected subset.

2.2. Categorization with respect to search strategies

To balance the trade-off of result optimality and computational efficiency, different search strategies such as exhaustive, complete, heuristic, and random search have been studied to generate candidate feature subsets for evaluation.

Exhaustive Search. An exhaustive approach to FS problem requires examining all possible subsets of size $d < D$ of the set X_D and selecting the subset with the largest value of J . Although the exhaustive search guarantees the optimality of the solution according to the evaluation criterion used, it is computationally prohibitive for problems of high dimensionality.

Complete Search. The alternative to exhaustive search is the *Branch & Bound* (B&B) algorithm, proposed first in [19], and ancestor algorithms based on a similar principle. All B&B algorithms rely on the *monotonicity property* of the FS criterion: given two subsets of the feature set X_D , A and B , if $A \subset B$, then $J(A) < J(B)$. By a straightforward application of this property many feature subset evaluations may be omitted. Therefore these algorithms guarantee to select an optimal feature subset of size d without involving explicit evaluation of all the possible combinations of d features. This assumption precludes the use of the error rate as the criterion in the classification problem. This is an important drawback as the error rate can be considered superior to other criteria [27], [10]. The time complexity is exponential in terms of data dimensionality, therefore the B&B algorithm is still prohibitive in terms of computation requirements for several recent applications, particularly those in data mining and document classification involving hundreds to tens of thousands of features.

Since searching for an optimal subset requires prohibitive computation, many sub-optimal search methods have been proposed, namely heuristic (also known as sequential) and random searching methods.

Heuristic Search. A number of heuristic (non-exhaustive) selection methods have been proposed which essentially trade-off the guarantee of optimality of the selected subset for computational complexity. This category includes, e.g., Best Individual Features, Sequential Forward and Backward Selection, Plus- l -Take Away- r , their

generalized versions, Genetic algorithms, and particularly the Floating and Oscillating Search. A comprehensive list of sub-optimal procedures can be found, e.g., in books [3], [5], [34], [9] [34]. A comparative taxonomy can be found, e.g., in [4], [8] or [11]. The time complexity of these methods is quadratic or less in terms of data dimensionality.

Random Search. It starts with a feature set containing all randomly selected features, add or remove features randomly selected one at a time and stop after a given number of iterations. The time complexity can be linear to the number of iterations. Various heuristic methods attempt to take use of the random search principle, hoping to avoid local extremes in the search space.

2.3. Categorization with respect to selection criteria

Based on the type of the *selection criterion*, feature selection algorithms may roughly be divided into three types: the *filter* [37], [2], the *wrapper* [10] and the *hybrid* [1], [26].

Filter methods are based on performance evaluation functions calculated directly from the training data such as *distance*, *information*, *dependency*, and *consistency*, and select feature subsets without involving any learning algorithm.

Wrapper approach requires one predetermined learning algorithm and uses its performance as the evaluation criterion. It attempts to find features better suited to the learning algorithm aiming to improve its performance. In the wrapper category, *classifier error rate* is used for classification and *cluster goodness* for clustering. Generally, the wrapper method achieves better performance than the filter method, but tends to be more computationally expensive than the filter approach. Also, the wrappers yield feature subsets optimized for the given learning algorithm only - the same subset may thus be bad in another context.

Hybrid approach combines the advantages of both approaches. Hybrid algorithms have recently been proposed to deal with high dimensional data. In these algorithms, first, a goodness measure of feature subsets based on data characteristics is used to choose best subsets for a given cardinality, and then, cross validation is exploited to decide a final best subset across different cardinalities. These algorithms mainly focus on combining filter and wrapper algorithms to achieve best possible performance with a particular learning algorithm with the time complexity comparable to that of the filter algorithms.

Filters as all other FS methods may be divided into global and local type. The *global* evaluation of features takes into account all data in a context-free way. The *local* FS, for example for classification tasks, means that different feature subsets may be used for different classes.

3. ADVANCED BRANCH & BOUND TYPE ALGORITHMS

In the field of complete search we have proposed some of the currently best-performing methods. Before discussing the more advanced algorithms, let us briefly summarize the essential Branch & Bound principle. The algorithm constructs a search tree where the root represents the set of all D features, X_D , and leaves represent target

subsets of d features. While tracking the tree down to leaves the algorithm successively removes single features from the current set of “candidates” (\mathcal{X}_k in k -th level). The algorithm keeps the information about both the till-now best subset of cardinality d and the corresponding criterion value, denoted as the *bound*. Anytime the criterion value in some internal node is found to be lower than the current *bound*, due to the required monotonicity property of the adopted criterion function the whole sub-tree may be cut-off and many computations may be omitted. The course of the B&B algorithm can be seen in Fig.1. (the meaning of C, P and A_i symbols is not important for now). This scheme in its simplest form is known as the “Basic B&B” algorithm. For details see [3], [5].

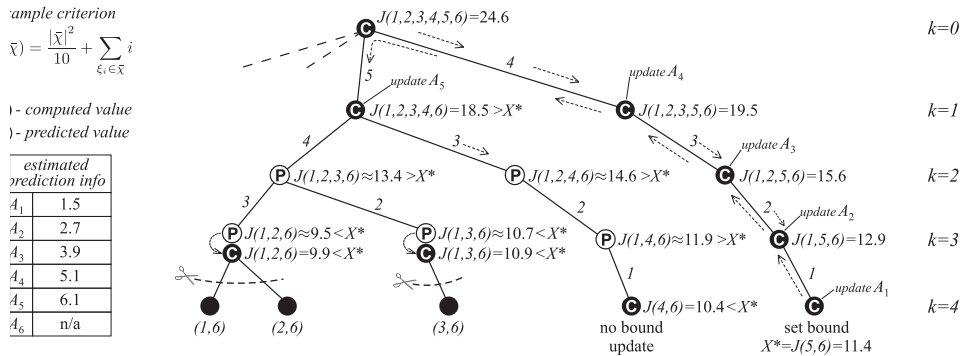


Fig. 1. Example of a “Fast Branch & Bound” problem solution, where $d = 2$ features are to be selected from a set of $D = 5$ features. Dashed arrows illustrate the way of tracking the search tree.

3.1. Branch & Bound properties

When compared to the exhaustive search, every B&B algorithm requires additional computations. Not only the target subsets of d features \bar{X}_{D-d} , but also their supersets \bar{X}_{D-d-j} , $j = 1, \dots, D - d$ have to be evaluated. The B&B principle does not guarantee enough sub-tree cut-offs to keep the total number of criterion computations lower than in exhaustive search.

To reduce the amount of criterion computations an additional node-ordering heuristic has been introduced in a more powerful “Improved B&B” (IBB) algorithm [3], [5]. IBB optimizes the order of features to be assigned to tree edges so that the *bound* value can increase as fast as possible and thus enables more effective branch cutting in later stages. Although IBB usually outperforms all simpler B&B algorithms, the computational cost of the additional heuristics can become a strong deteriorating factor. For detailed discussion of B&B drawbacks see [31]. In the following we present our recent effective framework for substantial B&B acceleration.

3.2. Prediction mechanism based Branch & Bound

The Fast Branch & Bound (FBB) [31] algorithm aims to reduce the number of criterion function computations in internal search tree nodes. A simplified algorithm description is as follows: FBB attempts to utilize the knowledge of past feature-dependent *criterion value decreases* (difference between criterion values before and after feature removal) for future prediction of criterion values without the need of real computation. Prediction is allowed under certain conditions only, e.g., not in leaves. Both the really computed and predicted criterion values are treated as equal while imitating the full IBB functionality, i.e., in ordering node descendants in the tree construction phase.

If the predicted criterion value remains significantly higher than the current *bound*, we may expect that even the actual value would not be lower and the corresponding sub-tree could not be cut-off. In this situation the algorithm continues to construct the consecutive tree level. However, if the predicted value is equal or lower than the *bound* (and therefore there arises a chance that the real value is lower than the *bound*), the real criterion value must be computed. Only if real criterion values are lower than the current *bound*, sub-trees may be cut-off. Note that this prediction scheme does not affect the optimality of obtained results. The FBB algorithm course remains similar to that of the IBB, possible sub-tree cut-offs are allowed according to real criterion values only. Possible inaccurate predictions may result in nothing worse than constructing sub-trees, which would have been pruned out by means of classical B&B algorithms. However, this situation is usually strongly outweighed by criterion computation savings in other internal nodes, especially near the root, where the criterion computation tends to be slower.

The prediction mechanism processes the information about the averaged *criterion value decrease* separately for each feature. The idea is illustrated in Fig. 1. For a detailed and formal description of this rather complex procedure and other B&B related topics see [31].

3.3. Improving the “Improved” algorithm

The FBB operates mostly the fastest among all B&B algorithms. However, it exhibits some drawbacks: it cannot be used with recursive criterion forms and there is no theoretical guarantee that extensive prediction failures won't hinder the overall speed, despite the fact that such faulty behavior has not been observed with real data. The B&B with Partial Prediction (BBPP) [31] constitutes a less effective but robust alternative. While learning similarly to FBB, it does not use predictions to substitute true criterion values inside the tree. Predicted values are used only for ordering before features get assigned to tree edges. In this sense BBPP can be looked upon as a slightly modified IBB with the only difference in node ordering heuristics. The performance gain follows from the fact, that the original IBB ordering heuristics always evaluates more criterion values than is the number of features finally used. For detailed analysis of BBPP see [31]. Among other recent B&B related ideas the “trading space for speed” approach [7] deserves attention as an alternative that may operate exceptionally fast under certain circumstances. The BBPP and FBB

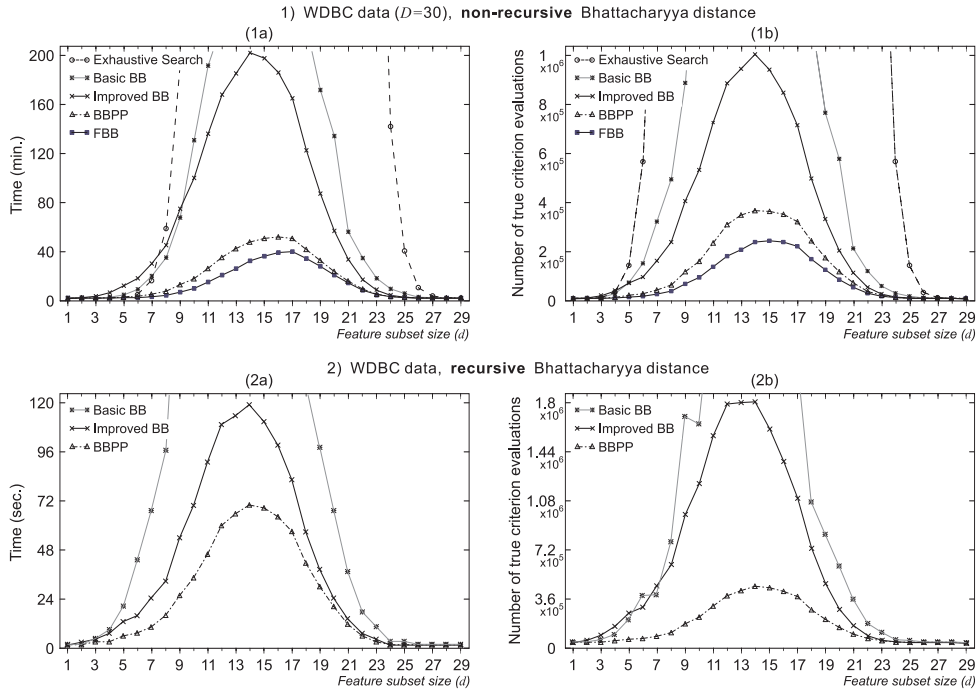


Fig. 2. Optimal subset search methods performance when maximizing the Bhattacharyya distance on 30-dimensional data (Wisconsin Diagnostic Breast Center).

algorithms are further investigated in [32], [35].

3.4. Predictive B&B properties and experimental results

When compared to classical B&B algorithms the predictive algorithms always spend additional time for maintaining the prediction mechanism. However, this additional time showed not to be a factor, especially when compared to time savings arising from the pruned criterion computations. The algorithms have been thoroughly tested on a number of different data sets. Here we show representative results on 30-dimensional mammogram data (2 classes – 357 benign and 212 malignant samples, see [18]). We used both the recursive (where applicable) and non-recursive Bhattacharyya distance as the criterion function. Performance of different methods is illustrated in Fig. 2.

We compare all results especially against the IBB algorithm [3], [5], as this algorithm has been long accepted to be the most effective optimal subset search method. Remark: Where applicable, we implement all algorithms to support “minimum solution tree” [36].

3.5. Summary of optimal methods

The only optimal subset search method usable with non-monotonic criteria is the exhaustive (full) search. However, because of exponential nature of the search problem, alternative methods are often needed. Our several recent improvements of the B&B idea especially in the form of prediction based FBB and BBPP resulted in a speed-up factor of 10 to 100 over the simplest B&B form, depending on particular data and criterion used. Among other recent B&B related ideas the “trading space for speed” approach [7] deserves attention as it may be capable of further improving the performance stated above under certain circumstances.

It should be stressed that despite the shown advances all optimal methods remain exponential in nature. If there is no particular need to request the global optimum of results, sub-optimal search methods offer greater flexibility and acceptable speed even for high-dimensional problems, while the solutions found are not necessarily much worse than the optimum. Note, that the optimality is with respect to the chosen criterion. Criterion choice may affect the final system performance more than the possibly slight difference between the globally and locally optimal feature selection.

4. HEURISTIC SEARCH

The simplest yet widely used *sequential forward (or backward) selection* methods [3], SFS (SBS), iteratively add (remove) one feature at a time so as to maximize the intermediate criterion value until the required dimensionality is achieved. Among the more interesting recent approaches the following two families of methods can be pointed out for general applicability and performance reasons, namely *sequential floating search methods* [22], [28] and *oscillating search methods* [29], both developed in UTIA.

Earlier sequential methods suffered from the so-called nesting of feature subsets which significantly deteriorated the performance. The first attempt to overcome this problem was to employ either the Plus- l -Take away- r [denoted (l, r)] or generalized (l, r) algorithms [3] which involve successive augmentation and depletion process. The same idea in a principally extended and refined form constitutes the basis of floating search.

4.1. Sequential floating search

The sequential forward floating selection procedure consists of applying after each forward step a number of backward steps as long as the resulting subsets are better than previously evaluated ones at that level. Consequently, there are no backward steps at all if intermediate result at actual level (of corresponding dimensionality) cannot be improved. The same applies for the backward version of the procedure. Both algorithms allow a ‘self-controlled backtracking’ so they can eventually find good solutions by adjusting the trade-off between forward and backward steps dynamically. In a certain way, they compute only what they need without any parameter setting. Formal description of this now classical procedure can be found in

[22].

Floating search algorithms can be considered the universal tools not only outperforming all predecessors, but also keeping advantages not met by more sophisticated algorithms. They find good solutions in all problem dimensions in one run. The overall search speed is high enough for most of ordinary problems.

4.2. Adaptive floating search

As the floating search algorithms have been found successful and generally accepted to be an efficient universal tool, their idea was further investigated. The so-called Adaptive Floating Search (AFS) has been proposed in [28]. The AFS algorithms are able to outperform the classical floating algorithms in certain cases, but at a cost of considerable increase of search time and the necessity to deal with unclear parameters. Our experience shows that AFS is usually inferior to newer algorithms, which we focus on in the following.

4.3. Oscillating search

The recent Oscillating Search (OS) [29] can be considered a “higher level” procedure, that takes use of other feature selection methods as sub-procedures in its own search. The concept is highly flexible and enables modifications for different purposes. It has shown to be very powerful and capable of over-performing standard sequential procedures, including Floating Search. Unlike other methods, the OS is based on repeated modification of the current subset X_d of d features. In this sense the OS is independent on the predominant search direction. This is achieved by alternating so-called *down-* and *up-swings*. Both *swings* attempt to improve the current set X_d by replacing some of the features by better ones. The *down-swing* first removes, then adds back, while the *up-swing* first adds, then removes. Two successive opposite swings form an *oscillation cycle*. The OS can thus be looked upon as a controlled sequence of oscillation cycles. The value of o (denoted *oscillation cycle depth*) determines the number of features to be replaced in one swing. o is increased after unsuccessful oscillation cycles and reset to 1 after each X_d improvement.

The algorithm terminates when o exceeds the user-specified *limit* Δ , or it may be left to reach the limits (full or empty set of features).

Every OS algorithm requires some initial set of d features. The initial set may be obtained randomly or in any other way, e.g., using some of the traditional sequential selection procedures. Furthermore, almost any feature selection procedure can be used in *up-* and *down-swings* to accomplish the replacements of feature o -tuples. Therefore, for the sake of generality in the following descriptions let us denote the adding / removing of a feature o -tuple by $\text{ADD}(o)$ / $\text{REMOVE}(o)$.

4.3.1. Oscillating search – formal algorithm description

Step 1: (*Initialization*) By means of any feature selection procedure (or randomly) determine the initial set X_d of d features. Let $c = 0$. Let $o = 1$.

Step 2: (*Down-swing*) By means of $\text{REMOVE}(o)$ remove such o -tuple from X_d

to get new set X_{d-o} so that $J(X_{d-o})$ is maximal. By means of $\text{ADD}(o)$ add such o -tuple from $X_D \setminus X_{d-o}$ to X_{d-o} to get new set X'_d so that $J(X'_d)$ is maximal. If $J(X'_d) > J(X_d)$, let $X_d = X'_d$, $c = 0$, $o = 1$ and go to Step 4.

Step 3: (*Last swing has not improved the solution*) Let $c = c + 1$. If $c = 2$, then nor the last *up*- nor *down-swing* led to a better solution. Extend the search by letting $o = o + 1$. If $o > \Delta$, stop the algorithm, otherwise let $c = 0$.

Step 4: (*Up-swing*) By means of $\text{ADD}(o)$ add such o -tuple from $X_D \setminus X_d$ to X_d to get new set X_{d+o} so that $J(X_{d+o})$ is maximal. By means of $\text{REMOVE}(o)$ remove such o -tuple from X_{d+o} to get new set X'_d so that $J(X'_d)$ is maximal. If $J(X'_d) > J(X_d)$, let $X_d = X'_d$, $c = 0$, $o = 1$ and go to Step 2.

Step 5: (*Last swing has not improved the solution*) Let $c = c + 1$. If $c = 2$, then nor the last *up*- nor *down-swing* led to a better solution. Extend the search by letting $o = o + 1$. If $o > \Delta$, stop the algorithm, otherwise let $c = 0$ and go to Step 2.

4.3.2. Oscillating search properties

The generality of the OS search concept allows to adjust the search for better speed or better accuracy (lower Δ and simpler ADD / REMOVE vs. higher Δ and more complex ADD / REMOVE). In this sense let us denote *sequential OS* the simplest possible OS version which uses a sequence of SFS steps in place of $\text{ADD}()$ and a sequence of SBS steps in place of $\text{REMOVE}()$. As opposed to all sequential search procedures, OS does not waste time evaluating subsets of cardinalities too different from the target one. The fastest improvement of the target subset may be expected in initial phases of the algorithm, because of the low initial cycle depth. Later, when the current feature subset evolves closer to optimum, low-depth cycles fail to improve and therefore the algorithm broadens the search (by letting $o = o + 1$). Though this improves the chance to get closer to optimum, the trade-off between finding a better solution and computational time becomes more apparent. Consequently, the OS tends to improve the solution most considerably during the fastest initial search stages. This behavior is advantageous, because it gives the option of stopping the search after a while without serious result-degrading consequences. Let us summarize the key OS advantages:

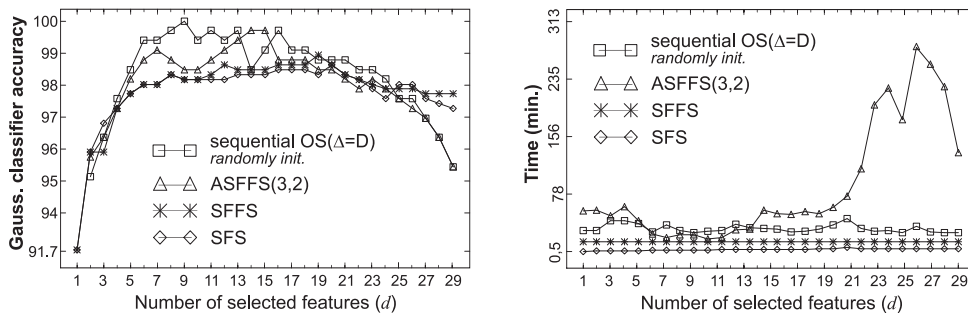


Fig. 3. Comparison of sub-optimal methods on classification problem.

- It may be looked upon as a universal tuning mechanism, being able to improve solutions obtained in another way.
- The randomly initialized OS is very fast, in case of very high-dimensional problems may become the only applicable procedure. E.g., in document analysis for search of the best 1000 words out of a vocabulary of 50000 even the simple SFS may show to be too slow.
- Because the OS processes subsets of target cardinality from the very beginning, it may find solutions even in cases, where the sequential procedures fail due to numerical problems.
- Because the solution improves gradually after each oscillation cycle, with the most notable improvements at the beginning, it is possible to terminate the algorithm prematurely after a specified amount of time to obtain a usable solution. The OS is thus suitable for use in real-time systems.
- In some cases the sequential search methods tend to uniformly get caught in certain local extremes. Running the OS from several different random initial points gives better chances to avoid that local extreme.

4.4. Hybrid floating search

We have defined a flexible hybrid version of floating search methods [33] for feature selection. The main benefit of the proposed floating search hybridization is the possibility to deal flexibly with the quality-of-result vs. computational time trade-off and to enable wrapper based feature selection in problems of higher dimensionality than before. We have shown that it is possible to trade significant reduction of search time for often negligible decrease of the classification accuracy.

4.5. Experimental results of sub-optimal search methods

All described sequential search methods have been tested on a large number of different problems. Here we demonstrate their performance on 2-class, 30-dimensional mammogram data (see [18]). The graphs in Fig. 3. show the OS ability to outperform other methods even in the simplest *sequential OS* form (here with $\Delta = d$ in one randomly initialized run). The ASFFS behavior is well illustrated here showing better performance than SFFS at a cost of uncontrollably increased time. SFFS and SFS need one run only to get all solutions. SFFS performance is always better than that of SFS.

4.6. Recommendations for practitioners

Concerning our current experience, we can give the following recommendations. Floating Search can be considered the tool of the first choice. It is reasonably fast and yields generally very good results in all dimensions at once, often succeeding in finding the global optimum. The Oscillating Search becomes better choice whenever: 1) the highest quality of solution must be achieved but optimal methods are

not applicable, or 2) a reasonable solution is to be found as quickly as possible, or 3) numerical problems hinder the use of sequential methods, or 4) extreme problem dimensionality prevents any use of sequential methods, or 5) the search is to be performed in real-time systems. Especially when applied repeatedly with different random starting points the Oscillating Search shows outstanding potential to overcome local extremes in favor of global optimum.

It should be stressed that, as opposed to B&B, the Floating Search and Oscillating Search methods are tolerant to deviations from monotonic behavior of feature selection criteria. It makes them particularly useful in conjunction with non-monotonic FS criteria like the error rate of a classifier (cf. Wrappers [10]), which according to a number of researchers seem to be the only legitimate criterion for feature subset evaluation.

Note: Floating and Oscillating Search source codes in C can be downloaded from <http://ro.utia.cas.cz/dem.html>.

5. MIXTURE BASED FEATURE SELECTION FOR CLASSIFICATION

Following the statistical approach to pattern recognition, we assume that a pattern represented by a feature vector $\mathbf{x} = (x_1, \dots, x_D)^T$ in D -dimensional feature space is to be classified into one of a finite set of C different classes $\{\omega_1, \dots, \omega_C\}$. A pattern \mathbf{x} belonging to class ω_j is viewed as an observation of a random vector drawn randomly according to the known class-conditional probability density function (pdf) $p(\mathbf{x}|\omega_j)$ and the respective *a priori* probability $P(\omega_j)$. We can say, that a tested pattern \mathbf{x} can be optimally classified using Bayes classification rule based on the knowledge of $P(c_j)p(\mathbf{x}|\omega_j)$ for each class ω_j as follows:

assign \mathbf{x} to class ω_l if

$$\omega_l = \arg \max_{j=1, \dots, C} \frac{P(\omega_j)p(\mathbf{x}|\omega_j)}{p(\mathbf{x})}, \quad l \in \{1, \dots, C\}, \quad p(\mathbf{x}) = \sum_{j=1}^C P(\omega_j)p(\mathbf{x}|\omega_j), \quad (2)$$

where $p(\mathbf{x})$ is unconditional pdf of \mathbf{x} .

In practice, there appear situations, where the class-conditional pdf is unknown and only a training set of labeled patterns is available. Finding a proper pdf estimate has a crucial impact to successful classification. Often simple models, such as a single Gaussian distribution, can effectively represent patterns but a more general model, such as a finite mixture model, must be used to approximate more complex pdfs; arbitrarily complex pdfs can be approximated using finite mixture models. For details on mixture models see e.g. [17].

In our papers [24], [23], [20], we proposed a mixture based classification with global feature selection by casting it as an estimation problem, thus avoiding any combinatorial search. Instead of selecting a subset of features, we estimate a set of binary quantities, one for each feature. This estimation is carried out by an Expectation-Maximization (EM) algorithm derived for the task.

5.1. Finite mixture model

It is said that D -dimensional random vector \mathbf{x} follows in class ω_j a finite mixture distribution with M_j components if its pdf can be written as

$$p(\mathbf{x}|\omega_j) = \sum_{m=1}^{M_j} \alpha_m^j p_m(\mathbf{x}|\theta_m^j), \quad (3)$$

where α_m^j is the *mixing probability* for the m th mixture component within the class ω_j , ($\alpha_m^j \geq 0, m = 1, \dots, M_j, \sum_{m=1}^{M_j} \alpha_m^j = 1$), θ_m^j is the set of parameters defining the m -th component, and $\{\theta_1^j, \dots, \theta_{M_j}^j, \alpha_1^j, \dots, \alpha_{M_j}^j\}$ is the complete set of parameters specifying the mixture model (3). This model assumes that each mixture component has a multivariate distribution with its own parameters.

A simplification of the model (3) consists of assuming that given the mixture component label, the features are conditionally statistically independent. In this case the j th class conditional pdf can be rewritten as

$$p(\mathbf{x}|\omega_j) = \sum_{m=1}^{M_j} \alpha_m^j \prod_{i=1}^D f(x_i|\theta_{mi}^j), \quad (4)$$

where $f(x_i|\theta_{mi}^j)$ is the pdf of the feature x_i in the m th mixture component within the class ω_j ; θ_{mi}^j is the set of parameters of the component m corresponding to feature x_i . Therefore, the class-conditional densities are modeled as a mixture of independent probabilistic models. This model has been considered e.g. by McLachlan in [16].

5.2. Mixture model with feature significance

In this section we introduce the concept of *feature significance* for the mixture structure of the class conditional pdf.

Assume now that some features are not significant for the mixture structure of the pdf defined in (4), in the following sense: if feature x_i is *not significant*, then its distribution is common to all mixture components. It means that any specific univariate function $f(x_i|\theta_{mi}^j)$ in the model (4) is substituted by the "background" density, denoted by $f_0(x_i|\theta_{0i})$, whenever the feature x_i is not significant. Let $\Phi = (\phi_1, \dots, \phi_D) \in \{0, 1\}^D$ be a set of binary parameters, such that $\phi_i = 1$ if the feature x_i is *significant* and $\phi_i = 0$, otherwise. In this way the binary parameters ϕ_i can be looked upon as *control variables* due to that the complexity and the structure of the mixture (4) can be controlled by means of that parameters. Then the class conditional pdf defined in (4) can be rewritten as

$$p(\mathbf{x}|\alpha_j, \theta_j, \theta_0, \Phi) = \sum_{m=1}^{M_j} \alpha_m^j \prod_{i=1}^D [f_0(x_i|\theta_{0i})^{1-\phi_i} f(x_i|\theta_{mi}^j)^{\phi_i}] \quad (5)$$

$$\alpha_j = (\alpha_1^j, \dots, \alpha_{M_j}^j), \quad \theta_j = (\theta_1^j, \dots, \theta_{M_j}^j), \quad \theta_0 = (\theta_{01}, \dots, \theta_{0D}), \quad \Phi = (\phi_1, \dots, \phi_D).$$

The novel characteristic of the mixture model defined in (5) are the *control variables* ϕ_i . When $\phi_i = 1$, it means that the feature x_i is significant and should be modeled by the individual mixture components. If $\phi_i = 0$, the feature x_i is useless and should be represented by a common component.

5.3. A feature significance measures for classification

In order to optimize the accuracy of the Bayes classification rule (2), the EM algorithm is used to find the unknown parameters $\{\alpha_j, \theta_j, \theta_0, \Phi\}$, $j = 1, \dots, C$ of the model (5) under the following assumptions:

- the training set of labeled samples is available;
- the univariate density functions f and f_0 belong to the family of univariate Gaussian densities;
- density function f_0 is common to all the classes ω_j , $j = 1, \dots, C$.

The criterion we use for measuring the error resulting from approximating the true probability density function $p(\mathbf{x}|\omega_j)$ by $p(\mathbf{x}|\alpha_j, \theta_j, \theta_0, \Phi)$ for all $j = 1, \dots, C$ is a mixture, in the true proportions $P(\omega_1), \dots, P(\omega_C)$, of the Kullback-Leibler distances between the true and the postulated class conditional pdf of \mathbf{x} .

Given the approximations of class conditional pdf's it can be easily seen that the "background" pdf f_0 may be reduced in the inequality in the Bayes plug-in classification rule and we may classify the tested pattern \mathbf{x} into one of C classes according to only d features x_{i_1}, \dots, x_{i_d} , where d is the required number of the selected features and $\{i_1, \dots, i_d\}$ is the permutation of $\{1, \dots, D\}$. We call this classifier in the reduced feature space the pseudo-Bayes classifier.

The proposed model (5) is suitable especially if the underlying distributions of some or all given classes are multi-modal or they conceal the existence of sub-populations. It allows to develop two FS methods: 'approximation' method determines those features which produce a set of approximations to the class probability density functions which is best in the sense of minimizing the mixture of Kullback-Leibler distances; 'divergence' method identifies those features that are most useful, in the sense of maximizing the the mixture of Kullback J-divergence, in describing differences between two possible classes.

Our approach to FS has the following features:

- it yields the feature subset of required size, optimizing the used criterion, without involving any search procedure;
- it provides a pseudo-Bayes plug-in classification rule employing the selected features.

Law et al. [12] (and references therein) has been inspired by the main idea of our approach to FS. They proposed a solution to the FS problem in unsupervised learning.

6. FEATURE SELECTION TOOLBOX

The Feature Selection Toolbox (FST) software [30] has been serving as a platform for data testing, feature selection, approximation-based modelling of data, classification and mostly testing newly developed methods. It is used basically for pattern

| | | gauss | approx. 1c | approx. 5c | approx. 10c | approx. 20c |
|---------------|------------------------|-------|---------------|---------------|----------------|----------------|
| <i>speech</i> | (random init.) | 8.39 | 21.61 | 7.58 | 9.19 | 9.03 |
| <i>data</i> | (dogs & rabbits init.) | - | 21.61 | 7.42 | 6.45 | 8.39 |
| <i>mammo</i> | (random init.) | 5.96 | 5.26 | 5.26 | 5.96 | 4.56 |
| <i>data</i> | (dogs & rabbits init.) | - | 5.26 | 5.26 | 5.96 | 5.96 |

Table 1. Error rates [%] of different classifiers with different parameters. The 'gauss' column holds results of a gaussian classifier. Other columns hold results obtained using the 'approximation' method (in this case the 'divergence' method yielded the same results). Note: 5c means 5 components of mixture, etc.

recognition purposes. However, we used it for solving decision making problems in economics and in other application fields as well.

A rather simple user interface was constructed upon a strong functional kernel. Most of results are generated in the form of textual protocol into the Console window. Numerical results may be collected in tables and used for generating graphs. Data may be displayed in a 2D projection.

6.1. Mixture-based FS for classification task example

We used the following real data-sets:

- 2-class, 15-dimensional *speech* data representing words "yes" and "no" obtained from the British Telecom; classes are separable with great difficulty.
- 2-class, 30-dimensional *mammogram* data representing benign and malignant patients, obtained from the Wisconsin Diagnostic Breast Center via the UCI repository [18].

Using the FST we compared the performance of gaussian classifier to the pseudo-Bayes classifier, defined especially for use with multimodal data, and defined in relation to "approximation" and "divergence" methods (c.f. section 5.). Table 1 illustrates the potential of the approximation model based classifiers. However, it also illustrates the necessity of experimenting to find a suitable number of components (the issue is discussed, e.g., in Sardo [25]).

The results were computed on the full set of features. In case of the "approximation" and "divergence" methods the algorithms were initialized randomly (1st row) by means of the "dogs & rabbits" cluster analysis (McKenzie et al. [15]) pre-processor (2nd row). Classifiers were trained on the first half of the dataset and tested on the second half.

Table 1. demonstrates the potential of mixture approximation methods – with 5 mixture components (see column approx.5c) for the "speech" data and 1, 5 or 20 components for the "mammo" data. The underlying data structure has been modeled precisely enough to achieve better classification rate when compared to the gaussian classifier. Second row for each data contains approximation and divergence

method results after preliminary initialization by means of the “dogs and rabbits” clustering method.

7. CONCLUSIONS AND FUTURE WORK

The current state of art in feature selection based dimensionality reduction for modeling and decision problems of classification type have been over-viewed. A number of recent feature subset search strategies have been reviewed and compared with emphasis put on methods developed within UTIA. Recent developments of B&B based algorithms for optimal search led to considerable improvements of the speed of search. Nevertheless, the principal exponential nature of optimal search remains and will remain one of key factors motivating the development of sub-optimal strategies. Among the family of sequential search algorithms the Floating and Oscillating search methods deserve particular attention. Two alternative feature selection methods based on mixture modelling have been presented. They are suitable for cases, when no *a priori* information on underlying probability structures is known. Among the most recent developments the hybrid FS methods deserve particular attention. This is reflected in our current field of interest [33]. In the future we intend to “hybridize” other search methods in a similar way and to investigate in detail the hybrid behavior of different combinations of various probabilistic measures and learning methods.

Many of recent feature selection methods have been implemented in Feature Selection Toolbox. The software has been used to demonstrate the differences between different criteria and differently selected feature subsets. The importance of feature selection for classification performance has been clearly shown as well.

ACKNOWLEDGEMENT

This work has been supported by the following grants: CR MŠMT grant 2C06019 and 1M0572 DAR, EC project FP6-507752 MUSCLE, Grant Agency of the Academy of Sciences of the Czech Republic (CR) No. A2075302.

(Received Month day, 2004.)

REFERENCES

- [1] S. Das: Filters, Wrappers and a boosting-based hybrid for Feature Selection. In: Proc. 18th Intl Conf. Machine Learning, 2001, pp. 74–81
- [2] M. Dash, K. Choi, P. Scheuermann, H. Liu: Feature selection for clustering - a Filter solution. In: Proc. Second Int. Conf. Data Mining, 2002, pp. 15–122
- [3] P.A. Devijver and J. Kittler: Pattern Recognition: A Statistical Approach. Prentice-Hall, 1982.
- [4] F.J. Ferri, P. Pudil, M. Hatef, J. Kittler: Comparative study of techniques for large-scale feature selection. In: Pattern Recognition in Practice IV, (E.S. Gelsema and L.N. Kanal), Elsevier Science B.V., 1994, pp. 403–413.
- [5] K. Fukunaga: Introduction to Statistical Pattern Recognition. Academic Press, 1990.
- [6] T. Hastie and R. Tibshirani: Discriminant analysis by Gaussian mixtures. Journal Royal Statist. Soc. B, 58, (1996), 155–176.

- [7] X. Chen: An improved Branch and Bound algorithm for feature selection. *Pattern Recognition Letters*, *24*(12), (2003), 1925–1933.
- [8] A.K. Jain and D. Zongker: Feature selection: Evaluation, application and small sample performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *19*(2), (1997), 153–158.
- [9] A.K. Jain, R.P.W. Duin, and J. Mao: Statistical Pattern Recognition: A Review, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, *22*(2): 4–37, 2000.
- [10] R. Kohavi and G.H. John: Wrappers for feature subset selection. *Artificial Intelligence*, *97*(1-2), (1997), 273–324.
- [11] M. Kudo and J. Sklansky: Comparison of algorithms that select features for pattern classifiers. *Pattern Recognition*, *33*(1), (2000), 25–41.
- [12] M.H. Law, M.A.T. Figueiredo and A.K. Jain: Simultaneous feature selection and clustering using mixture models. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, *26*, (2004), 1154–1166.
- [13] H. Liu and L. Yu: Toward integrating feature selection algorithms for classification and clustering. *IEEE Trans. on Knowledge and Data Engineering*, *17*, (2005), 491–502.
- [14] H.A. Mayer, P. Somol, R. Huber and P. Pudil: Improving statistical measures of feature subsets by conventional and evolutionary approaches. In: *Proc. 3rd IAPR International Workshop on Statistical Techniques in Pattern Recognition (Alicante)*, 2000, pp. 77–81.
- [15] P. McKenzie and M. Alder: Initializing the EM algorithm for use in gaussian mixture modelling. University of Western Australia, 1994.
- [16] G.J. McLachlan: *Discriminant Analysis and Statistical Pattern Recognition*. John Wiley & Sons, New York 1992.
- [17] G.J. McLachlan and D. Peel: *Finite Mixture Models*. John Wiley & Sons, New York 2000.
- [18] P.M. Murphy and D.W. Aha: UCI Repository of Machine Learning Databases [ftp.ics.uci.edu]. University of California, Department of Information and Computer Science, Irvine, CA, 1994.
- [19] P.M. Narendra and K. Fukunaga: A branch and bound algorithm for feature subset selection. *IEEE Trans. on Computers*, *26*, (1977), 917–922.
- [20] J. Novovičová, P. Pudil and J. Kittler: Divergence based feature selection for multimodal class densities. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, *18*(2), (1996), 218–223.
- [21] J. Novovičová and P. Pudil: Feature selection and classification by modified model with latent structure. In: *Dealing With Complexity: Neural Network Approach*, Springer Verlag, 1997, pp. 126–140.
- [22] P. Pudil, J. Novovičová and J. Kittler: Floating search methods in feature selection. *Pattern Recognition Letters*, *15*(11), (1994), 1119–1125.
- [23] P. Pudil, J. Novovičová and J. Kittler: Feature selection based on approximation of class densities by finite mixtures of special type. *Pattern Recognition*, *28*, (1995), 1389–1398.
- [24] P. Pudil, J. Novovičová and J. Kittler: Simultaneous learning of decision rules and important attributes for classification problems in image analysis. *Image and Vision Computing*, *12*, (1994), 193–198.
- [25] L. Sardo and J. Kittler: Model complexity validation for PDF estimation using Gaussian mixtures. In: *Proc. 14th International Conference on Pattern Recognition*, 2, 1998, pp. 195–197.
- [26] M. Sebban and R. Nock: A Hybrid filter/wrapper approach of feature selection using information theory. *Pattern Recognition*, *35*, (2002), 835–846.

- [27] W. Siedlecki and J. Sklansky: On Automatic Feature Selection. *International Journal of Pattern Recognition and Artificial Intelligence*, 2(2), (1988), 197–220.
- [28] P. Somol, P. Pudil, J. Novovičová and P. Paclík: Adaptive floating search methods in feature selection. *Pattern Recognition Letters*, 20(11,12,13), (1999), 1157–1163.
- [29] P. Somol and P. Pudil: Oscillating search algorithms for feature selection. In: *Proc. 15th IAPR International Conference on Pattern Recognition*, 2000, pp. 406–409.
- [30] P. Somol and P. Pudil: Feature Selection Toolbox. *Pattern Recognition*, 35(12), (2002), 2749–2759.
- [31] P. Somol, P. Pudil and J. Kittler: Fast branch & bound algorithms for optimal feature selection. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 26(7), (2004), 900–912.
- [32] P. Somol, P. Pudil and J. Grim: On prediction mechanisms in fast branch & bound algorithms. In: *Lecture Notes in Computer Science 3138*, Springer, Berlin 2004, pp. 716–724.
- [33] P. Somol, J. Novovičová and P. Pudil: Flexible-Hybrid Sequential Floating Search in Statistical Feature Selection. In: *Lecture Notes in Computer Science*, Springer, 2006 (to appear).
- [34] S. Theodoridis and K. Koutroumbas: *Pattern Recognition*. 2nd Ed. Academic Press, 2003.
- [35] Z. Wang, J. Yang and G. Li: An improved branch & bound algorithm in feature selection. In: *Lecture Notes in Computer Science LNCS 2639*, Springer, 2003, pp. 549–556.
- [34] A. Webb: *Statistical Pattern Recognition*. 2nd Ed. John Wiley & Sons, 2002.
- [36] B. Yu and B. Yuan: A more efficient branch and bound algorithm for feature selection. *Pattern Recognition*, 26, (1993), 883–889.
- [37] L. Yu and H. Liu: Feature selection for high-dimensional data: A fast correlation-based filter solution. In: *Proc. 20th Intl Conf. Machine Learning*, 2003, pp. 856–863.

Petr Somol, Institute of Information Theory and Automation – Academy of Sciences of the Czech Republic, Pod vodárenskou věží 4, 182 08 Praha 8 and Faculty of Management – Prague University of Economics, Czech Republic.

Jana Novovičová, Institute of Information Theory and Automation – Academy of Sciences of the Czech Republic, Pod vodárenskou věží 4, 182 08 Praha 8 and Faculty of Management – Prague University of Economics Czech Republic. Czech Republic.

Pavel Pudil, Faculty of Management – Prague University of Economics, 377 01 Jindřichův Hradec, Czech Republic, and Institute of Information Theory and Automation – Academy of Sciences of the Czech Republic, Pod vodárenskou věží 4, 182 08 Praha 8.

e-mails: novovic, somol@utia.cas.cz pudil@fm.vse.cz