

Feature Subset Selection by Bayesian networks: a comparison with genetic and sequential algorithms

¹Department of Computer Science and Artificial Intelligence.

University of the Basque Country.

P.O. Box 649.

E-20080 Donostia - San Sebastián.

Basque Country.

Spain.

Telf: (+34) 943015026.

Fax: (+34) 943219306.

e-mail: ccbincal@si.ehu.es

Abstract

In this paper we perform a comparison among FSS-EBNA, a randomized, population-based and evolutionary algorithm, and two genetic and other two sequential search approaches in the well known Feature Subset Selection (FSS) problem. In FSS-EBNA, the FSS problem, stated as a search problem, uses the EBNA (Estimation of Bayesian Network Algorithm) search engine, an algorithm within the EDA (Estimation of Distribution Algorithm) approach. The EDA paradigm is born from the roots of the GA community in order to explicitly discover the relationships among the features of the problem and not disrupt them by genetic recombination operators. The EDA paradigm avoids the use of recombination operators and it guarantees the evolution of the population of solutions and the discovery of these relationships by the factorization of the probability distribution of best individuals in each generation of the search. In EBNA, this factorization is carried out by a Bayesian network induced by a cheap local search mechanism. FSS-EBNA can be seen as an hybrid Soft Computing system, a synergistic combination of probabilistic and evolutionary computing to solve the FSS task. Promising results on a set of real Data Mining domains are achieved by FSS-EBNA in the comparison respect to well known genetic and sequential search algorithms.

Keywords: Feature Subset Selection, Estimation of Distribution Algorithm, Soft Computing, Estimation of Bayesian Network Algorithm, Bayesian network, predictive accuracy

I. MOTIVATION

The basic problem of Supervised Classification in Data Mining is concerned with the induction of a model that classifies a given object into one of several known classes. In order to induce the classification model, each object is described by a pattern of d features. With advanced computer technologies, big data archives are usually formed and many features are used to describe the objects. Here, Data Mining and Machine Learning communities usually formulate the following question: *‘Are all of these d descriptive features useful for learning the classification model?’* On trying to respond to this question, we come up with the Feature Subset Selection (FSS) approach, which can be reformulated as follows: *given a set of candidate features, select the best subset in a classification task.*

The dimensionality reduction made by a FSS process can carry out several advantages for a classification system in a specific task:

- a reduction in the cost of acquisition of the data;
- improvement of the comprehensibility of the final classification model;
- a faster induction of the final classification model;
- an improvement in classification accuracy.

The attainment of higher classification accuracies, coupled with a notable dimensionality reduction, is the common objective of Machine Learning and Data Mining processes.

It has long been proved that the classification accuracy of supervised classifiers is not monotonic with respect to the addition of features. Irrelevant or redundant features, depending on the specific characteristics of the supervised classifier, may degrade the predictive accuracy of the classification model.

FSS can be viewed as a search problem, with each state in the search space specifying a subset of the possible features of the task. Exhaustive evaluation of possible feature subsets is usually unfeasible in practice because of the large amount of computational effort required. Due to its randomized, evolutionary and population-based nature, Genetic Algorithms (GAs) have been the most commonly used search engine in the FSS process (Kudo and Sklansky [32], Siedelecky and Sklansky [54], Vafaie and De Jong [58]). Most of the theory of GAs deals with the so-called Building Blocks (BBs, Goldberg [18]); simply said, by BBs, partial solutions of a problem are meant, formed by groups of related variables. GAs reproduce BBs by an implicit manipulation of a large number of them by mechanisms of selection and

recombination. A crucial factor of the GA success resides in the proper growth and mixing of the optimal BBs of the problem. Problem-independent recombination operators often break these BBs and do not mix them efficiently; thus, this could delay the discovery of the global optima or produce a convergence to a local optima.

Linkage learning (LL) (Harik and Goldberg [20]) is the identification of the BBs to be conserved under recombination. Recently, various approaches to solve the LL problem have been proposed. Several proposed methods are based on the manipulation of the representation of solutions during the optimization to make the interacting components of partial solutions less likely to be broken. For this purpose, various reordering and mapping operators are used. One of these approaches is the well known Messy Genetic Algorithm (MGA) (Goldberg [18]). Instead of extending the GA, in latter years, a new approach has strongly emerged under the EDA term (Estimation of Distribution Algorithm, Mühlenbein and Paaß[43], Larrañaga and Lozano [37]). EDA approach explicitly learns the probabilistic structure of the problem and uses this information to ensure a proper mixing and growth of BBs that do not disrupt them. The further exploration of the search space is guided, instead of crossover and mutation operators as in GAs, by the probabilistic modeling of promising solutions.

Bearing these difficulties in mind, in this paper we perform a comparison in the FSS task among FSS-EBNA [25], a novel algorithm inspired on the EDA approach, and other classic genetic and sequential search algorithms. FSS-EBNA (Feature Subset Selection by Estimation of Bayesian Network Algorithm) is based on the probabilistic modeling by Bayesian networks (BNs) of promising solutions (feature subsets, in our case) to guide further exploration of the space of features. This BN is sampled to obtain the next generation of feature subsets, therefore, guaranteeing the evolution of the search for the optimal subset of features. Using a wrapper approach to evaluate the goodness of each candidate feature subset, we search for the feature subset which maximizes the predictive accuracy of the Naive-Bayes supervised classifier on presented tasks. However, this paper presents the first comparison among FSS-EBNA and other classic search algorithms in the FSS task.

FSS-EBNA can be seen as an hybrid Soft Computing (Bonissone et al. [7]) system, a synergistic combination of probabilistic and evolutionary computing to solve the FSS task: the use of a classic probabilistic paradigm as Bayesian networks in conjunction with the evolutionary strategy scheme of the EDA paradigm, enables us to construct an application which tackles the FSS problem. The remainder of the paper is organized as follows: the next sec-

tion introduces the FSS problem and its basic components. Section 3 introduces the EDA approach, BNs and the combination of both, the EBNA search engine. Section 4 presents the application of EBNA to the FSS problem, resulting in the FSS-EBNA algorithm. In section 5, we compare it against four well known FSS methods, specially focusing our attention on the comparison against two genetic approaches. The last section picks up the summary of our work and some ideas about future lines of research.

II. FEATURE SUBSET SELECTION AS A SEARCH PROBLEM

Even if we locate our work as a Machine Learning or Data Mining approach, the FSS literature includes plenty of works in other fields such as Pattern Recognition (Jain and Chandrasekaran [27], Kittler [29]), Statistic (Miller [40], Narendra and Fukunaga [46]) or Text-Learning (Mladenić [41]). In the Bayesian network community we have the example of the work of Provan and Singh [51], who, using a Bayesian network as a classifier, build it selecting in a greedy manner the variables that should maximize the predictive accuracy of the network. Thus, different communities have exchanged and shared ideas among them to deal with the FSS problem.

As reported by Aha and Bankert [2], the objective of FSS in Machine Learning and Data Mining is to *reduce the number of features used to characterize a dataset so as to improve a classifier's performance on a given task*. In our framework we will appreciate a tradeoff between a high-predictive-accuracy and a low-cardinality (selection of few features) in the finally chosen feature subset. The FSS task can be exposed as a search problem, each state in the search space identifying a subset of possible features.

Thus, any FSS algorithm must determine four basic issues that define the nature of the search process (see Liu and Motoda [39] for a review of FSS algorithms):

1. *The starting point in the space*. It determines the direction of the search. One might start with no features and successively add them, or one might start with all features and successively remove them. One might also select an initial state somewhere in the middle of the search space.
2. *The organization of the search*. It determines the strategy of the search. Roughly speaking, the search strategies can be *complete* or *heuristic*. The basis of the *complete search* is the systematic examination of every possible feature subset. Three classic complete search implementations are depth-first, breadth-first and branch and bound search (Narendra and

Fukunaga [46]). On the other hand, among *heuristic* algorithms, there are *deterministic heuristic* algorithms and *non-deterministic heuristic* ones. Classic deterministic heuristic FSS algorithms are sequential forward selection and sequential backward elimination (Kittler [29]), floating selection methods (Pudil et al. [52]) and best-first search (Kohavi and John [30]). They are deterministic in the sense that all the runs always obtain the same solution. *Non-deterministic heuristic* search is used to escape from local maximums. Randomness is used for this purpose and this implies that one should not expect the same solution from different runs. Two classic implementations of non-deterministic search engines are the frequently applied Genetic Algorithms (Siedelecky and Sklansky [54]) and Simulated Annealing (Doak [15]).

3. *Evaluation strategy of feature subsets.* By the calculation of the goodness of each proposed feature subset the evaluation function identifies the promising areas of the search space. The objective of FSS algorithm is its maximization. The search algorithm uses the value returned by the evaluation function to guide the search. Some evaluation functions carry out this objective looking only at the intrinsic characteristics of the data and measuring the power of a feature subset to discriminate among the classes of the problem: this type of evaluation functions are grouped below the *filter* strategy. These evaluation functions are usually monotonic and increase with the addition of features that afterwards can hurt the predictive accuracy of the final classifier. However, Kohavi and John [30] reported that when the goal of FSS is the maximization of the accuracy, the features selected should depend not only on the features and the target concept to be learned, but also on the special characteristics of the supervised classifier. Thus, they proposed the *wrapper* concept: this implies that the FSS algorithm conducts the search for a good subset using the classifier itself as a part of the evaluation function, the same classifier that will be used to induce the final classification model. Once the classification algorithm is fixed, the idea is to train it with the feature subset found by the search algorithm, estimating the predictive accuracy on the training set and assigning it as the value of the evaluation function of the feature subset.

4. *Criterion for halting the search.* An intuitive approach for stopping the search will be the non-improvement of the evaluation function value of alternative subsets. Another classic criterion will be to fix an amount of possible solutions to be visited along the search.

In this section, EDA paradigm and BNs will be explained. Bearing these two concepts in mind, EBNA, the search engine used in our FSS algorithm will be presented. EDA paradigm is the general formula of the EBNA algorithm; on the other hand, Bayesian networks can be seen as the most important basis of EBNA.

A. The need of the EDA paradigm

Many combinatorial optimization algorithms have no mechanism for capturing the relationships among the variables of the problem. The related literature has many papers proposing different heuristics in order to implicitly capture these relationships. GAs implicitly capture these relationships by concentrating samples on combinations of high-performance members of the current population through the use of the crossover operator. Crossover combines the information contained within pairs of selected ‘parent’ solutions by placing random subsets of each parent’s bits into their respective positions in a new ‘child’ solution. In GAs no explicit information is kept about which groups of variables jointly contribute to the quality of candidate solutions. The crossover operation is randomized and could disrupt many of these relationships among the variables; therefore, most of the crossover operations yield unproductive results and the discovery of the global optima could be delayed.

On the other hand, GAs are also criticized in the literature for three aspects [36]:

- the large number of parameters and their associated referred optimal selection or tuning process;
- the extremely difficult prediction of the movements of the populations in the search space;
- their incapacity to solve the well known deceptive problems [18].

Linkage Learning (LL) is the identification of groups of variables (or Building Blocks, BBs) that are related. Instead of extending the GA, the idea of the explicit discovery of these relationships during the optimization process itself has emerged from the roots of the GA community. One way to discover these relationships is to estimate the joint distribution of promising solutions and use this estimate in order to generate new individuals¹. A general scheme of the algorithms based on this principle is called the Estimation of Distribution Algorithm (EDA, Mühlenbein and Paß[43]). In EDA (see Figure 1), there are no crossover

¹In a search algorithm the terms ‘individual’ and ‘solution’ can be used indistinctly.

-
1. $D_0 \leftarrow$ Generate N individuals (the initial population) randomly.
 2. Repeat for $l = 1, 2, \dots$ until a stop criterion is met:
 - 2.1. $D_{l-1}^s \leftarrow$ Select $S \leq N$ individuals from D_{l-1} according to a selection method.
 - 2.2. $p_l(\mathbf{x}) = p(\mathbf{x}|D_{l-1}^s) \leftarrow$ Estimate the joint probability distribution of an individual being among the selected individuals.
 - 2.3. $D_l \leftarrow$ Sample N individuals (the new population) from $p_l(\mathbf{x})$.
-

Fig. 1. Main scheme of the EDA paradigm.

nor mutation operators, and the new population is sampled from a probability distribution which is estimated from the selected individuals. However, estimating the distribution is a critical task in EDA (for a more complete review see Larrañaga et al. [36]).

The simplest way to estimate the distribution of good solutions assumes the independence between the features² of the problem. New candidate solutions are sampled by only regarding the proportions of the values of the variables independently to the remaining ones. Population Based Incremental Learning (PBIL, Baluja [5]), Compact Genetic Algorithm (cGA, Harik et al. [20]), Univariate Marginal Distribution Algorithm (UMDA, Mühlenbein [42]) and Bit-Based Simulated Crossover (BSC, Syswerda [57]) are four algorithms of this type. They work well in problems with no significant interactions (or relationships) among features and so, the need for covering higher order interactions among the variables is seen for more complex or real tasks.

Efforts covering pairwise interactions among the features of the problem have generated algorithms such as: MIMIC uses simple chain distributions (De Bonet et al. [14]), the so-called dependency trees (Baluja and Davies [6]) and the Bivariate Marginal Distribution Algorithm (BMMA, Pelikan and Mühlenbein [50]). However, Pelikan and Mühlenbein [50] have demonstrated that only covering pairwise dependencies is not enough with problems which have higher order interactions: in this type of problems the discovery of the global optimum is considerably delayed when lower order based probability estimating approaches are used.

In this way, the Factorized Distribution Algorithm (FDA, Mühlenbein and Mahnig [44]) covers higher order interactions. This is done using a previously fixed factorization of the

²In the Evolutionary Computation community, the term ‘variable’ is normally used instead of ‘feature’. We use both terms indistinctly.

joint probability distribution. However, FDA needs prior information on the decomposition and factorization of the problem which is often not available. Without the need of this extra information about the decomposition and factorization of the problem, other probabilistic models, able to cover higher order interactions, appear in the literature. For this purpose, Bayesian networks, graphical representations able to cover these higher order interactions, can be used. Thus, EBNA (Etxeberria and Larrañaga [16]) BOA (Pelikan et al. [49]) are algorithms which use different implementations of Bayesian networks for estimating the joint distribution of promising solutions. In this way, multivariate interactions among variables can be covered. In Soto et al. [56] a Bayesian network with a polytree structure is proposed: the proposed algorithm is called PADA (Polytree Approximation of Distribution Algorithms) and it is based on a method proposed by Acid and de Campos [1] for detecting relationships among variables. Ochoa et al. [47] propose an initial algorithm to learn (from conditional independence tests) the junction tree from a database. Harik presents an algorithm (Extend compact Genetic Algorithm, EcGA [21]), whose basic idea consists of factorizing the joint probability distribution in a product of marginal distributions of variable size: these marginal distributions of variable size are related to the variables that are contained in the same group and to the probability distributions associated to them.

Many of these works notify of a faster discovery of the global optimum by EDA algorithms than GAs for certain combinatorial optimization problems. Harik [21] and Pelikan and Mühlenbein [50] show several intuitive and well known problems where the GAs, when the crossover operator is applied, frequently disrupt optimum relationships among features: in this way, these optimum relationships that appear in the parent solutions will disappear in the children solutions and the discovery of the global optimum will be delayed. Thus, these authors note that EDA approaches are able to first discover and then maintain these relationships during the entire optimization process, producing a faster discovery of the global optimum than GAs.

In order to avoid the evaluation of larger amounts of possible solutions (and its associated CPU time), a faster discovery of well fitted areas in the search space is highly desired in the FSS problem. Bearing this purpose in mind, we use an EDA inspired approach for FSS. Based on the EBNA work of Etxeberria and Larrañaga [16], we propose the use of Bayesian networks as the models for representing the probability distribution of the set of candidate solutions for the FSS problem, using the application of automatic learning methods to induce

the right distribution model in each generation in an efficient way.

B. Bayesian networks

A Bayesian network (BN, Castillo et al. [9], Lauritzen [38], Pearl [48]) encodes the relationships contained in the modelled data. It can be used to describe the data as well as to generate new instances of the variables with similar properties as those of given data. A Bayesian network encodes the probability distribution $p(\mathbf{x})$, where $\mathbf{X} = (X_1, \dots, X_d)$ is a vector of variables, and it can be seen as a pair (M, θ) . M is a directed acyclic graph (DAG) where the nodes correspond to the variables and the arcs represent the conditional (in)dependencies³ (Dawid [13]) among the variables. By X_i , both the variable and the node corresponding to this variable is denoted. M will give the factorization of $p(\mathbf{x})$:

$$p(\mathbf{x}) = \prod_{i=1}^d p(x_i | \pi_i)$$

where Π_i is the set of parent variables (nodes) that X_i has in M and π_i its possible instantiations. The number of states of Π_i will be denoted as $|\Pi_i| = q_i$ and the number of different values of X_i as $|X_i| = r_i$. $\theta = \{\theta_{ijk}\}$ are the required conditional probability values to completely define the joint probability distribution of \mathbf{X} . θ_{ijk} will represent the probability of X_i being in its k th state while Π_i is in its j th instantiation. This factorization of the joint distribution can be used to generate new instances using the conditional probabilities in a modeled dataset.

Informally, an arc between two nodes relates these nodes so that the value of the variable corresponding to the ending node of the arc depends on the value of the variable corresponding to the starting node. Every probability distribution can be defined by a Bayesian network (Pearl [48]). As a result, Bayesian networks are widely used in problems where uncertainty is handled using probabilities.

Unfortunately, finding the optimal \hat{M} requires searching through all possible structures, which has been proven to be NP-hard (Chickering et al. [11]). Although promising results have been obtained using global search techniques (Chickering et al. [12], Larrañaga et al. [34], Larrañaga et al. [35], Wong et al. [59]) their computation cost makes them unfeasible for our problem. We need to find \hat{M} as fast as possible, so a simple algorithm which returns a good structure, even if it is not optimal, will be preferred.

³Let \mathbf{Y}, \mathbf{Z} and \mathbf{W} be three disjoint sets of variables, then \mathbf{Y} is said to be conditionally independent of \mathbf{Z} given \mathbf{W} , if and only if $p(\mathbf{y} | \mathbf{z}, \mathbf{w}) = p(\mathbf{y} | \mathbf{w})$, for all possible values \mathbf{y}, \mathbf{z} and \mathbf{w} of \mathbf{Y}, \mathbf{Z} and \mathbf{W} .

In our implementation Algorithm B (Buntine [8]) is used for learning Bayesian networks from data. Algorithm B is a greedy search heuristic. The algorithm starts with an arc-less structure in each generation of the search and at each step, it adds the arc with the maximum increase in the BIC approximation (Penalized maximum likelihood of the proposed structure, Schwarz [55]). The algorithm stops when adding an arc does not increase the utilized measure. In the BIC metric, the penalized function is given by the Jeffreys–Schwarz criterion, resulting in the Bayesian Information Criterion (BIC [55]). The BIC score of a Bayesian network structure M , from a database D and containing N cases (characterised by d variables) is denoted as follows:

$$BIC(M, D) = \sum_{i=1}^d \sum_{j=1}^{q_i} \sum_{k=1}^{r_i} N_{ijk} \log \frac{N_{ijk}}{N_{ij}} - \frac{\log N}{2} \sum_{i=1}^d (r_i - 1) q_i$$

N_{ijk} denotes the number of cases in D in which the variable X_i is instantiated to its k -th value and its parents in the network structure are instantiated to their j -th value. Let be $N_{ij} = \sum_{k=1}^{r_i} N_{ijk}$, where r_i and q_i represent the number of possible states of X_i and its parents respectively.

C. Estimation of Bayesian Network Algorithm: EBNA

The general procedure of EBNA appears in Figure 2. To understand the steps of the algorithm, the following concepts must be clarified.

M_0 is the DAG with no arcs at all and $\theta_0 = \{\forall i : p(X_i = x_i) = \frac{1}{r_i}\}$, which means that the initial Bayesian network BN_0 , assigns the same probability to all individuals. N is the number of individuals in the population. S is the number of individuals selected from the population. Although S can be any value, the most common value in the EDA literature is $S = \frac{N}{2}$. If S is close to N then the populations will not evolve very much from generation to generation. On the other hand, a low S value will lead to low diversity resulting in early convergence.

For individual selection range based selection is proposed, i.e., selecting the best $N/2$ individuals from the N individuals of the population. However, any selection method could be used. Probabilistic Logic Sampling algorithm (PLS, Henrion [22]) is used to sample new individuals from the Bayesian network.

Finally, the way in which the new population is created must be pointed out. In the given procedure, all individuals from the previous population are discarded and the new popula-

-
1. $BN_0 \leftarrow (M_0, \theta_0)$.
 2. $D_0 \leftarrow$ Sample N individuals from BN_0 .
 3. For $l = 1, 2, \dots$ until a stop criterion is met:
 - 3.1. $D_{l-1}^s \leftarrow$ Select S individuals from D_{l-1} .
 - 3.2. $\hat{M}_l \leftarrow$ Find the structure which maximizes $BIC(M_l, D_{l-1}^s)$.
 - 3.3. $\theta_l \leftarrow$ Calculate $\{\theta_{ijk}^l = \frac{N_{ijk}^{l-1} + 1}{N_{ij}^{l-1} + r_i}\}$ using D_{l-1}^s as the data set.
 - 3.4. $BN_l \leftarrow (\hat{M}_l, \theta_0)$.
 - 3.5. $D_l \leftarrow$ Sample N individuals from BN_l using PLS.
-

Fig. 2. EBNA basic scheme.

tion is composed of all newly created individuals⁴. This has the problem of losing the best individuals that have been previously generated, therefore, the following minor change has been made: instead of discarding all the individuals, we maintain the best individual of the previous generation and create $N - 1$ new individuals.

An elitist approach is used to form iterative populations. Instead of directly discarding the $N - 1$ individuals from the previous generation replacing them with $N - 1$ newly generated ones, the $2N - 1$ individuals are put together and the best $N - 1$ taken among them. These best $N - 1$ individuals form the new population together with the best individual of the previous generation. In this way, the populations converge faster to the best individuals found; however, this also implies a risk of losing diversity within the population.

The stopping criteria will be explained in the next section.

IV. FEATURE SUBSET SELECTION BY ESTIMATION OF BAYESIAN NETWORK

ALGORITHM: FSS-EBNA. INSTANTANATION OF FSS COMPONENTS

Once the FSS problem and EBNA algorithm are presented, we will use the search engine provided by EBNA to solve the FSS problem, resulting in the FSS-EBNA algorithm. In the next lines we will briefly describe its basic components (see Inza et al. [25] for more details). Parametrizing the FSS task as an optimization problem, it consists on selecting, for a training set characterised by d features, the subset of features which maximizes the accuracy of a certain classifier on further test instances. Thus, the cardinality of the search space is 2^d . Using an intuitive notation to represent each individual (there are d bits in each individual,

⁴Newly created individuals are also known as ‘offspring’.

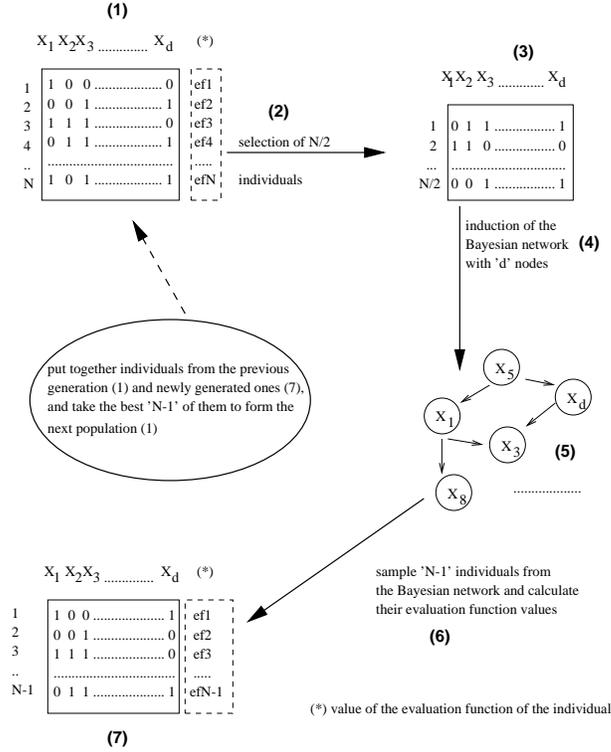


Fig. 3. FSS-EBNA method.

each bit indicating whether a feature is present (1) or absent (0)), we can see in Figure IV an overview of the application of the EBNA search engine to the FSS problem. In each generation of the search, the induced Bayesian network will factorize the probability distribution of selected individuals. The Bayesian network will be formed by d nodes, each one representing one feature of the domain.

The determination of a minimum population size to reliably estimate the parameters of a Bayesian network is not an easy task (Friedman and Yakhini [17]). This difficulty is higher for real world problems where the true real probability distribution is not known. In this way, taking the dimensionality of our problems into account, we consider that a population size of 1,000 individuals is enough to reliably estimate the Bayesian network parameters.

A wrapper approach will be used to calculate the evaluation function value of each proposed individual or feature subset. The value of the evaluation function of a feature subset found by the EBNA search technique, once the supervised classifier is fixed, will be calculated by an accuracy estimation on training data.

The accuracy estimation, seen as a random variable, has an intrinsic uncertainty. A 10-fold cross-validation multiple times, combined with a heuristic proposed by Kohavi and John

[30], will be used to control the intrinsic uncertainty of the evaluation function. The heuristic works as follows:

- if the standard deviation of the accuracy estimate is above 1%, another 10-fold cross-validation is executed;
- this is repeated until the standard deviation drops below 1%, a maximum of five times.

In this way, small datasets will be cross-validated many times. However, larger ones possibly once.

Although FSS-EBNA is independent from the specific supervised classifier used within its wrapper approach, in our set of experiments we will use the well known Naive-Bayes (NB, Cestnik [10]) supervised classifier. It is a simple and fast classifier which uses the Bayes rule to predict the class for each test instance, assuming that features are independent to each other given the class. Due to its simplicity and fast induction, it is commonly used on Data Mining tasks of high dimensionality (Kohavi and John [30], Mladenić [41]). The probability for discrete features is estimated from data using maximum likelihood estimation and applying the Laplace correction. A normal distribution is assumed to estimate the class conditional probabilities for continuous attributes. Unknown values in the test instance are skipped. Although its simplicity and its independence assumption among variables, the literature shows that the NB classifier gives remarkably high accuracies in many domains (Langley and Sage [33]), specially in medical ones. Despite its good scaling with irrelevant features, NB can improve its accuracy level discarding correlated or redundant features. NB, based on the independence assumption of features to predict the class, is hurt by correlated features which violate this independence assumption. Thus, FSS can also play a ‘normalization’ role to discard these groups of correlated features, ideally selecting one of them in the final model. Although Langley and Sage [33] propose a sequential forward feature selection for detecting these correlations (starting from an empty subset of features and sequentially selecting features until no improvement is achieved), Kohavi and John [30] prefer a sequential backward elimination (starting from the full set of features and sequentially deleting features until no improvement is achieved). In our work we propose the use of FSS-EBNA to detect these redundancies among features to improve NB’s predictive accuracy.

To stop the search algorithm, we have adopted an intuitive stopping criteria which takes the number of instances of the training set into account. In this way, we try to avoid the

‘overfitting’ risk (Jain and Zongker [28], Kohavi and John [30]):

- for datasets with more than 2,000 instances, the search is stopped when in a sampled new generation no feature subset appears with an evaluation function value improving the best subset found in the previous generation. Thus, the best subset of the search, found in the previous generation, is returned as the FSS-EBNA’s solution;
- for smaller datasets the search is stopped when in a sampled new generation no feature subset appears with an evaluation function value improving, at least with a p-value smaller than 0.1^5 , the value of the evaluation function of the best feature subset of the previous generation. Thus, the best subset of the previous generation is returned as the FSS-EBNA’s solution.

For larger datasets the ‘overfitting’ phenomenon has a lesser impact and we hypothesize that an improvement in the accuracy estimation over the training set will be coupled with an improvement in generalization accuracy on unseen instances. Otherwise, for smaller datasets, in order to avoid the ‘overfitting’ risk, the continuation of the search is only allowed when a significant improvement in the accuracy estimation of best individuals of consecutive generations appears. We hypothesize that when this significant improvement appears, the ‘overfitting’ risk decays and there is a basis for further generalization accuracy improvement over unseen instances.

Another concept to understand this stopping criteria is the wrapper nature of the proposed evaluation function. As we will see in the next section, the evaluation function value of each visited solution (the accuracy estimation of the NB classifier on training set by 10-fold cross-validation multiple times, using only the features proposed by the solution) needs several seconds to be calculated (never more than 4 seconds for used datasets). As the generation of a new generation of individuals implies the evaluation of 1,000 new individuals, we only allow the continuation of the search when it demonstrates that it is able to escape from the local optima, being able to discover new ‘best’ solutions in each generation. When the wrapper approach is used, the CPU time must be also controlled: we hypothesize that when the continuation of the search is allowed by our stopping criteria, the CPU times to evaluate a new generation of solutions are justified. For a larger study about this stopping criterion, the work [25] can be consulted.

⁵Using a 10-fold cross-validated paired t test between the folds of both estimations, taking only the first run into account when 10-fold cross-validation is repeated multiple times.

Table I. Details of experimental domains. C = continuous. D = discrete.

domain	number of instances	number of classes	number of features and their types
<i>Ionosphere</i>	351	2	34 (34-C)
<i>Horse-colic</i>	368	2	22 (15-D, 7-C)
<i>Soybean-large</i>	683	19	35 (35-C)
<i>Anneal</i>	898	6	38 (6-D, 32-C)
<i>Image</i>	2,310	7	19 (19-C)

V. EMPIRICAL COMPARISON

We will test the power of FSS-EBNA in five real datasets. Table 1 reflects the principal characteristics of the datasets. All datasets come from the *UCI repository* (Murphy [45]) and they have been frequently used in the FSS literature. Due to their large dimensionality, an exhaustive search for the best feature subset is not computationally feasible, so a heuristic search must be performed.

In real datasets we do not know the optimal subset of features respect to the Naive-Bayes classifier. To test the power of FSS-EBNA, a comparison with the following well known FSS algorithms is carried out:

- Sequential Forward Selection (SFS) is a classic hill-climbing search algorithm (Kittler [29]) which starts from an empty subset of features and sequentially selects features until no improvement is achieved in the evaluation function value. It performs the major part of its search near the empty feature set;
- Sequential Backward Elimination (SBE) is another classic hill-climbing algorithm (Kittler [29]) which starts from the full set of features and sequentially deletes features until no improvement is achieved in the evaluation function value. It performs the major part of its search near the full feature set. Instead of working with a population of solutions, SFS and SBE try to optimize a single feature subset;
- GA with one-point crossover (GA-o);
- GA with uniform crossover (GA-u);
- FSS-EBNA.

For all the FSS algorithms, the wrapper evaluation function explained in the previous section is used. Although SFS and SBE algorithms stop deterministically, GA algorithms apply the same stopping criteria as FSS-EBNA.

Although the optimal selection of parameters is still an open problem on GAs (Grefenstatte

Table II. Accuracy percentages of the Naive-Bayes (NB) classifier on real datasets without feature selection and using exposed five FSS methods. The last row shows the average accuracy percentages for six domains.

domain	without FSS	SFS	SBE	GA-o	GA-u	FSS-EBNA
<i>Ionosphere</i>	84.84 ± 3.12†	90.25 ± 1.58*	91.39 ± 2.68	91.17 ± 3.19	90.97 ± 2.56*	92.40 ± 2.04
<i>Horse-colic</i>	78.97 ± 2.98†	83.31 ± 1.98	82.12 ± 2.41*	83.43 ± 2.82	83.51 ± 1.47	83.93 ± 1.58
<i>Soybean-large</i>	81.96 ± 3.46†	86.38 ± 3.30*	87.78 ± 3.90*	85.64 ± 4.06†	86.09 ± 4.37†	88.64 ± 1.70
<i>Anneal</i>	93.01 ± 3.13*	86.72 ± 2.09†	92.49 ± 2.94*	92.95 ± 2.67*	93.13 ± 2.56	94.10 ± 3.00
<i>Image</i>	79.95 ± 1.52†	88.65 ± 1.21	88.82 ± 1.74	88.67 ± 2.48	89.12 ± 1.56	88.98 ± 0.98
Average	83.74	87.06	88.52	88.37	88.56	89.53

[19]), for both GA algorithms, guided by the recommendations of Bäck [4], the probability of crossover is set to 1.0 and the mutation probability to $1/d$, being d the number of variables of the domain (these values are so common in the literature). Fitness-proportionate selection (Goldberg [18]) is used to select individuals for crossover. In order to avoid any bias in the comparison, the remaining GA parameters are the same as FSS-EBNA’s: the population size is set to 1,000 and the new population formed by the best members from both old population and offspring.

Due to the non-deterministic nature of FSS-EBNA and GAs (two executions could not give the same result), 5 replications of 2-fold cross-validation (5x2cv) are applied to assess the predictive generalization accuracy of compared FSS algorithms. In each replication, the available data are randomly partitioned into two equal-sized sets S_1 and S_2 . The FSS algorithm is trained on each set and tested on the other set. In this way, the reported accuracies are the mean of ten accuracies. The standard deviation of the mean is also reported. We lengthen the comparison by running the Naive-Bayes classifier without feature selection. Table 2 shows the accuracy results on tested datasets. Apart from a high accuracy level, we will also focus our attention on the achievement of a reduced number of features: a good tradeoff between a high-accuracy and a low-cardinality of the selected feature subset is desired.

A deeper analysis of the accuracy results is carried out by means of statistical tests. The 5x2cv F (Alpaydin [3]) test is performed to determine the significance degree of accuracy differences among each approach and FSS-EBNA. Thus, in Table 2 the symbol ‘†’ denotes statistically significant difference over FSS-EBNA at the 0.05 confidence level; ‘*’, significance at the 0.1 confidence level. Table 3 shows the average (and its standard deviation) number of features selected by different approaches. Experiments are executed in a SGI-Origin 200 computer using the implementation of the Naive-Bayes algorithm that appears in the *MLC++* [31]

Table III. Cardinalities of finally selected features subsets for the Naive-Bayes (NB) classifier on real datasets without feature selection and using exposed five FSS methods. It must be taken into account that when no FSS is applied to NB, it uses all the features.

domain	without FSS	SFS	SBE	GA-o	GA-u	FSS-EBNA
<i>Ionosphere</i>	34	6.00 ± 1.41	21.30 ± 3.80	15.00 ± 2.36	12.66 ± 1.03	13.40 ± 2.11
<i>Horse-colic</i>	22	6.00 ± 2.74	11.20 ± 2.65	5.00 ± 2.82	4.60 ± 1.75	6.10 ± 1.85
<i>Soybean-large</i>	35	12.70 ± 2.71	23.50 ± 2.75	19.00 ± 2.09	19.16 ± 2.31	18.90 ± 2.76
<i>Anneal</i>	38	5.50 ± 2.32	33.60 ± 2.91	21.66 ± 2.66	19.50 ± 2.25	20.50 ± 3.13
<i>Image</i>	19	5.60 ± 1.57	9.40 ± 1.95	8.00 ± 1.41	8.00 ± 1.09	8.00 ± 0.66

software.

Regarding these tables, the following observations can be made:

- all FSS algorithms help NB to reduce the number of features needed to induce the final models. This dimensionality reduction is coupled with considerable accuracy improvements in all datasets except *Anneal*, where FSS-EBNA is the only algorithm able to significantly improve the accuracy of the NB model without feature selection;
- FSS-EBNA has the best average accuracy among compared algorithms. It has the best estimated value in all except one dataset (*Image* domain, where differences are not significant respect to the best algorithm, GA-u);
- by average SBE, GA-o and GA-u achieve similar accuracy results (their average accuracy results are also very similar). Their accuracies show significant differences respect to FSS-EBNA in *Soybean-large* (SBE, GA-o and GA-u), *Anneal* (SBE and GA-o), *Ionosphere* (GA-u) and *Horse-colic* (SBE);
- three randomized approaches (GA-o, GA-u, FSS-EBNA) carry out similar dimensionality reductions. We also note that in major part, they select the same features: thus, their accuracy differences in several datasets arise from their disjoint features;
- although SFS returns the subsets with the fewest features, its accuracy values are in all except one dataset significantly inferior to FSS-EBNA. As SFS starts from the empty feature set, it is biased to perform the major part of its search near the empty feature set (Doak [15]) and as we may expect it to return subsets with few features. However, it seems that it gets trapped on local optimas and it can not escape to better fitted areas;
- apart from the exposed accuracy differences between SBE and FSS-EBNA in several datasets, the major disadvantage of SBE is the slight reduction in the number of features it performs. In all domains, SBE is the algorithm with the lowest feature reduction, being

Table IV. Mean stop-generation for GAs and FSS-EBNA. The standard deviation of the mean is also reported. The initial generation is considered as the zero generation.

domain	GA-o	GA-u	FSS-EBNA
<i>Ionosphere</i>	3.50 ± 0.84†	3.10 ± 0.56†	1.80 ± 0.42
<i>Horse-colic</i>	3.20 ± 1.13*	3.40 ± 0.51†	2.40 ± 0.69
<i>Soybean-large</i>	3.30 ± 0.82*	3.60 ± 0.51†	2.50 ± 0.70
<i>Anneal</i>	3.80 ± 0.42†	3.20 ± 0.44†	1.80 ± 0.42
<i>Image</i>	3.60 ± 0.84	3.70 ± 0.48	3.50 ± 0.42

nearly insignificant in *Anneal* dataset. As SBE starts from the full feature subset, it is biased to perform the major part of its search near the full feature set (Doak [15]) and as we may expect it to return subsets with many features. However, it seems that it is not able to jump to better (or equally) fitted areas with fewer features;

- FSS-EBNA and GA approaches are not initially biased to specific areas of the search space and they are able to discover these well-fitted areas of non specific dimensionality. Vafaie and De Jong [58] demonstrate on a set of artificial and real problems how greedy-like sequential searches as SFS and SBE have a tendency to get trapped on local peaks of the search space, and they defend GAs as a more robust search engine in the FSS problem.

Although we have seen that two GA approaches and FSS-EBNA obtain similar accuracy results in many datasets, further studies show that GA approaches need more generations than FSS-EBNA to arrive to similar (or lower) accuracy levels. Table 4 reflects in which generation GA approaches and FSS-EBNA stop using the explained stopping criteria. The 5x2cv F test is performed to determine the significance degree of the differences in the stop-generation among each GA approach and FSS-EBNA. Thus, in Table 4 the symbol ‘†’ denotes a statistically significant difference over FSS-EBNA at the 0.05 confidence level; ‘*’, significance at the 0.1 confidence level.

The results show that FSS-EBNA arrives faster to similar or better predictive accuracies (see also Table 2) than both GA approaches. GA-o and GA-u show significant differences in the stop-generation respect to FSS-EBNA in all except one dataset. FSS-EBNA, by the use of Bayesian networks, captures the underlying structure of the problem faster than GAs, avoiding the disruption, as GAs, of the relationships among dependent variables. Only in the *Image* dataset, the domain of lowest dimensionality, the use of FSS-EBNA does not return an advantage respect to other algorithms: observing the accuracy and stop-generation results

in this domain, all FSS algorithms demonstrate a similar behaviour, finding similarly-fitted subsets and population-based approaches needing nearly the same number of generations to stop.

This faster discovery of optimal solutions by EDA approaches than GAs was also noted by Harik [21] and Pelikan and Muehlenbein [50] in a set of classical optimization problems. As in our case, when a wrapper approach is used to calculate the evaluation function value of each found subset, the faster discovery of similar or better accuracies is a critical task. Although NB is a fast supervised classifier, it needs several seconds to estimate the predictive accuracy (by the exposed 10-fold cross-validation multiple times) of a feature subset on the training set: depending on the number of features selected by the subset, near 1 CPU second is needed in *Ionosphere* (the domain with fewest instances), while near 3 are needed in *Image* (the domain with most instances). As the generation of a new population of solutions implies the evaluation of new 1,000 individuals, an earlier stop without a damage to the accuracy level is highly desired in order to save CPU time. On the other hand, the times for the induction of the Bayesian networks over the best individuals are insignificant in all domains: 3 CPU seconds are needed by average in *Image* domain (the domain with fewer features) and 14 CPU seconds in *Anneal* (the domain with more features).

In a subsequent analysis between GA-o and GA-u predictive accuracies, we note that the differences are not statistically significant in any dataset. So, we should focus our attention on their stop-generation numbers to extract differences between the behaviour of GA-o and GA-u. Thus, observing Table 4, the one-point crossover is better suited for *Horse-colic* and *Soybean-large* datasets than the uniform crossover; otherwise, we see the opposite behaviour in *Ionosphere* and *Anneal*. By the use of FSS-EBNA, we also avoid this tuning among different crossover operators for each dataset.

VI. SUMMARY AND FUTURE WORK

In this work, we have performed a comparison among FSS-EBNA, a randomized and population-based evolutionary algorithm, and two genetic and other two sequential classic search algorithms in the FSS task. In FSS-EBNA, the FSS problem, stated as a search problem, uses the EBNA (Estimation of Bayesian Network Algorithm) search engine, a variant of the EDA (Estimation of Distribution Algorithm) approach. The EDA paradigm is born from the roots of the GA community in order to explicitly discover the relationships among the features of

the problem and not disrupt them by genetic recombination operators. The EDA paradigm avoids the use of recombination operators and it guarantees the evolution of solutions and the discovery of these relationships by the factorization of the probability distribution of best individuals in each generation of the search. In EBNA, this factorization is carried out by a Bayesian network induced by a cheap local search mechanism. As another attractive effect, with the use of FSS-EBNA we avoid the tuning of genetic parameters, still an open problem on GA literature.

Using a wrapper approach, we have searched for the subset of features that optimizes the predictive accuracy of the Naive-Bayes supervised classifier. We have carried out a comparison in a set of tasks among two well known sequential search engines algorithms (Sequential Forward Selection and Sequential Backward Elimination) and Genetic Algorithms with one-point and uniform crossover operators. FSS-EBNA has demonstrated a robust behaviour, obtaining the best average accuracy respect to all the tasks, the best in all except one dataset and a large reduction in the number of selected features respect the full feature set. We have noted the tendency of Sequential Forward Selection to get trapped on local accuracy minima and the incapacity of Sequential Backward Elimination to perform an acceptable feature reduction. We have also noted the capacity of FSS-EBNA to arrive to similar or better fitted solutions faster than both GA approaches. Bayesian networks, which factorize the probability distribution of best solutions in each generation, are able to capture the underlying structure of the problems faster than GA approaches. On the other hand, the GA's crossover operator frequently disrupts the dependencies among related features: this is the cause of the slower discovery by GAs of solutions of the same level as FSS-EBNA.

As future work, we consider lengthening the works already done using the Bayesian networks for the Feature Weighting problem in Nearest Neighbor Algorithm (Inza [23]) and for the representation of the behaviour of Data Mining classifiers (Inza et al. [24]). Continuing the work within the EDA approach for FSS, in order to deal with domains with much larger numbers of features (> 50), future work should address the use of simpler probability models to factorize the probability distribution of best individuals, models which assume fewer or no dependencies between the variables of the problem (see Inza et al. [26] for a medical application). Other interesting possibility is the use of parallel algorithms to induce Bayesian networks in these tasks of high dimensionality (Sangüesa et al. [53], Xiang and Chu [60]). Another way of researching will be the employment of a metric which fixes for each domain the

number of individuals needed to reliably learn (Friedman and Yakhini [17]) the parameters of the Bayesian network.

Acknowledgements

This work is partially supported by the University of the Basque Country and by the Department of Education, University and Research of the Basque Government under grants 9/UPV/EHU/ 00140.226-12084/2000 and PI 1999-40 respectively.

REFERENCES

- [1] Acid, S., and de Campos, L.M., Approximations of causal networks by polytrees: an empirical study, in *Advances in Intelligent Computing, Lecture Notes in Computer Science 945*, (B. Bouchon-Meunier, R.R. Yager and L.A. Zadeh, Eds.), Springer-Verlag, Berlin, Germany, 149-158, 1995.
- [2] Aha, D.W., and Bankert, R.L., Feature selection for case-based classification of cloud types: An empirical comparison, *Proceedings of the AAAI'94 Workshop on Case-Based Reasoning*, 106-112, 1994.
- [3] Alpaydin, E., Combined 5x2cv F test for comparing supervised classification learning algorithms, *Neural Comput.*, 11, 1885-1982.
- [4] Bäck, T., *Evolutionary Algorithms is Theory and Practice*, Oxford University Press, 1996.
- [5] Baluja, S., Population-based incremental learning: A method for integrating genetic search based function optimization and competitive learning, Technical Report CMU-CS-94-163, Carnegie Mellon University, Pittsburgh, PA, 1994.
- [6] Baluja, S., and Davies, S., Using optimal dependency-trees for combinatorial optimization: Learning the structure of the search space, *Proceedings of the Fourteenth International Conference on Machine Learning*, 30-38, 1997.
- [7] Bonissone, P.P., Chen, Y.-T., Goebel, K., and Khedkar, P.S., Hybrid Soft Computing Systems: Industrial and Commercial Applications, *Proceedings of the IEEE*, 87(9), 1641-1667, 1999.
- [8] Buntine, W., Theory refinement in Bayesian networks, *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence*, 52-60, 1991.
- [9] Castillo, E., Gutiérrez, J.M., and Hadi, A.S., *Expert Systems and Probabilistic Network*

- Models*, Springer-Verlag, Berlin, Germany, 1997.
- [10] Cestnik, B., Estimating Probabilities: a crucial task in Machine Learning, *Proceedings of the European Conference on Artificial Intelligence*, 147-149, 1990.
 - [11] Chickering, M., Geiger, D., and Heckerman, D., Learning Bayesian networks is NP-hard, Technical Report MSR-TR-94-17, Microsoft Research, Redmond, WA, 1994.
 - [12] Chickering, D.M., Geiger, D., and Heckerman, D., Learning Bayesian networks: Search methods and experimental results, *Preliminary Papers of the 5th International Workshop on Artificial Intelligence and Statistics*, 112-128, 1995.
 - [13] Dawid, A.P., Conditional independence in statistical theory, *J. Roy. Stat. Soc. B*, 41, 1-31, 1979.
 - [14] De Bonet, J.S., Isbell, C.L., and Viola, P., MIMIC: Finding optima by estimating probability densities, in *Advances in Neural Information Processing Systems 9*, MIT Press, Cambridge, MA, 1997.
 - [15] Doak, J., An evaluation of feature selection methods and their application to computer security, Technical Report CSE-92-18, University of California at Davis, CA, 1992.
 - [16] Etxeberria, R., and Larrañaga, P., Global Optimization with Bayesian networks, *Proceedings of the Second Symposium on Artificial Intelligence*, La Habana, Cuba, 332-339, 1999.
 - [17] Friedman, N., and Yakhini, Z., On the Sample Complexity of Learning Bayesian Networks, *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence*, Portland, OR, 274-282, 1996.
 - [18] Goldberg, D.E., *Genetic algorithms in search, optimization, and machine learning*, Addison-Wesley, Reading, MA, 1989.
 - [19] Grefenstette, J.J., Optimization of Control Parameters for Genetic Algorithms, *IEEE T. Syst. Man. Cyb.*, 1, 122-128, 1986.
 - [20] Harik, G.R., Lobo, F.G., and Goldberg, D.E., The compact genetic algorithm, IlliGAL Report 97006, Urbana: University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, 1997.
 - [21] Harik, G., Linkage Learning via Probabilistic Modeling in the ECGA, IlliGAL Report 99010, Urbana: University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, 1999.
 - [22] Henrion, M., Propagating uncertainty in Bayesian networks by probabilistic logic sam-

- pling, in *Uncertainty in Artificial Intelligence 2*, Elsevier Science Publishers B.V., Amsterdam, The Netherlands, 149-163, 1988.
- [23] Inza, I., Feature Weighting for Nearest Neighbor Algorithm by Bayesian Networks based Combinatorial Optimization, *Proceedings of the Student Session of Advanced Course on Artificial Intelligence*, 33-35, 1999.
- [24] Inza, I., Larrañaga, P., Sierra, B., Etxeberria, R., Lozano, J.A., and Peña, J.M., Representing the behaviour of supervised classification learning algorithms by Bayesian networks, *Pattern Recogn. Lett.*, 20(11-13), 1201-1210, 1999.
- [25] Inza, I., Larrañaga, P., Etxeberria, R., and Sierra, B., Feature Subset Selection by Bayesian networks based optimization, *Artif. Intell.*, 123(1-2), 157-184, 2000.
- [26] Inza, I., Merino, M., Larrañaga, P., Quiroga, J., Sierra, B., and Giralá, M., Feature subset selection by genetic algorithms and estimation of distribution algorithms. A case study in the survival of cirrhotic patients treated with TIPS, *Artif. Intell. Med.* Accepted for publication.
- [27] Jain, A.K., and Chandrasekaran, R., Dimensionality and sample size considerations in pattern recognition practice, in *Handbook of statistics-II*, North-Holland, Amsterdam, The Netherlands, 835-855, 1982.
- [28] Jain, A., Zongker, D., Feature Selection: Evaluation, Application, and Small Sample Performance, *IEEE T. Pattern Anal.*, 19(2), 153-158, 1997.
- [29] Kittler, J., Feature Set Search Algorithms, *Pattern Recognition and Signal Processing*, Sithoff and Noordhoff, Alphen aan den Rijn, The Netherlands, 41-60, 1978.
- [30] Kohavi, R., and John, G., Wrappers for feature subset selection, *Artif. Intell.*, 97(1-2), 273-324, 1997.
- [31] Kohavi, R., Sommerfield, D., and Dougherty, J., Data mining using MLC++, a Machine Learning Library in C++, *International Journal of Artificial Intelligence Tools*, 6, 537-566, 1997.
- [32] Kudo, M., and Sklansky, J., Comparison of algorithms that select features for pattern classifiers, *Pattern Recogn.* 33, 25-41, 2000.
- [33] Langley, P., and Sage, S., Induction of selective Bayesian classifiers, *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence*, 399-406, 1994.
- [34] Larrañaga, P., Kuijpers, C.M.H., Murga, R.H., and Yurramendi, Y., Learning Bayesian network structures by searching for the best ordering with genetic algorithms, *IEEE T.*

- Syst. Man Cy. A*, 26(4), 487-493, 1996.
- [35] Larrañaga, P., Poza, M., Yurramendi, Y., Murga, R.H., and Kuijpers, C.M.H., Structure Learning of Bayesian networks by genetic algorithms: A performance analysis of control parameters, *IEEE T. Pattern Anal.*, 18(9), 912-926, 1996.
- [36] Larrañaga, P., Etxeberria, R., Lozano, J.A., and Peña, J.M., Combinatorial Optimization by Learning and Simulation of Bayesian Networks, in *Proceedings of the Conference in Uncertainty in Artificial Intelligence*, 343-352, 2000.
- [37] Larrañaga, P., and Lozano, J.A., *Estimation of Distribution Algorithms. A New Tool for Evolutionary Computation*, Kluwer Academic Publishers, Norwell, MA. To appear.
- [38] Lauritzen, S.L., *Graphical Models*, Oxford University Press, Oxford, England, 1996.
- [39] Liu, H., and Motoda, H., *Feature Selection for Knowledge Discovery and Data Mining*, Kluwer Academic Publishers, Norwell, MA, 1998.
- [40] Miller, A.J., *Subset Selection in Regression*, Chapman and Hall, Washington, DC, 1990.
- [41] Mladenić, D., Feature subset selection in text-learning, *Proceedings of the Tenth European Conference on Machine Learning*, 95-100, 1988.
- [42] H. Mühlenbein, 'The equation for response to selection and its use for prediction', *Evolutionary Computation*, 5(3), 303-346, (1997).
- [43] Mühlenbein, H., and Paaß, G., From recombination of genes to the estimation of distributions. Binary parameters, in *Lecture Notes in Computer Science 1411: Parallel Problem Solving from Nature - PPSN IV*, 178-187, 1996.
- [44] Mühlenbein, H., and Mahnig, T., FDA - A scalable evolutionary algorithm for the optimization of additively decomposed functions, *Evolutionary Computation*, 7(4), 353-376, 1999.
- [45] Murphy, P., UCI Repository of machine learning databases, Irvine, CA: University of California, Department of Information and Computer Science, 1995.
- [46] Narendra, P., and Fukunaga, K., A branch and bound algorithm for feature subset selection, *IEEE T. Comput.*, C-26(9), 917-922, 1977.
- [47] Ochoa, A., Soto, M., Santana, R., Madera, J., and Jorge, N., The Factorized Distribution Algorithm and the Junction Tree: A Learning Perspective, *Proceedings of the Second Symposium on Artificial Intelligence*, 368-377, 1999.
- [48] Pearl, J., *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, Palo Alto, CA, 1988.

- [49] Pelikan, M., Goldberg, D.E., and Cantú-Paz, E., *BOA: The Bayesian Optimization Algorithm*, IlliGAL Report 99003, Urbana: University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, 1999.
- [50] Pelikan, M., and Mühlenbein, H., The Bivariate Marginal Distribution Algorithm, in *Advances in Soft Computing-Engineering Design and Manufacturing*, Springer-Verlag, London, 521-535, 1999.
- [51] Provan, G.M., and Singh, M., Learning Bayesian networks using feature selection, *Preliminary Papers of the Fifth International Workshop on Artificial Intelligence and Statistics*, 450-456, 1995.
- [52] Pudil, P., Novovicova, J., and Kittler, J., Floating Search Methods in Feature Selection, *Pattern Recogn. Lett.*, 15(1), 1119-1125, 1994.
- [53] Sangüesa, R., Cortés, U., and Gisolfi, A., A parallel algorithm for building possibilistic causal networks, *Int. J. Approx. Reason.*, 18(3-4), 251-270, 1998.
- [54] Siedelecky, W., and Sklansky, J., On automatic feature selection, *Int. J. Pattern. Recogn.*, 2, 197-220, 1988.
- [55] Schwarz, G., Estimating the dimension of a model, *Ann. Stat.*, 7, 461-464, 1978.
- [56] Soto, M., Ochoa, A., Acid, S., and de Campos, L.M., Introducing Polytrees Approximation of Distribution Algorithm, *Proceedings of the Second Symposium on Artificial Intelligence*, 360-367, 1999.
- [57] Syswerda, G., Uniform Crossover in Genetic Algorithms, *Proceedings of the International Conference on Genetic Algorithms 3*, 2-9, 1989.
- [58] Vafaie, H., and De Jong, K., Robust feature selection algorithms, *Proceedings Fifth International Conference on Tools with Artificial Intelligence*, 356-363, 1993.
- [59] Wong, M.L., Lam, W., and Leung, K.S., Using Evolutionary Programming and Minimum Description Length Principle for Data Mining of Bayesian Networks, *IEEE T. Pattern Anal.* 21(2), 174-178, 1999.
- [60] Xiang, Y., and Chu, T., Parallel Learning of Belief Networks in Large and Difficult Domains, *Data Min. Knowl. Disc.* 3(3), 315-338, 1999.