Bioinformatics Advance Access published August 24, 2007

A review of feature selection techniques in bioinformatics

Yvan Saeys[1], Iñaki Inza[2] and Pedro Larrañaga[2]

[1] Department of Plant Systems Biology, VIB, B-9052 Ghent, Belgium and Bioinformatics and Evolutionary Genomics group, Department of Molecular Genetics, Ghent University, B-9052 Ghent, Belgium

[2] Department of Computer Science and Artificial Intelligence, Computer Science Faculty, University of the Basque Country, Paseo Manuel de Lardizabal 1, 20018 Donostia - San Sebastián, Spain

Associate Editor: Dr. Jonathan Wren

ABSTRACT

Feature selection techniques have become an apparent need in many bioinformatics applications. In addition to the large pool of techniques that have already been developed in the machine learning and data mining fields, specific applications in bioinformatics have led to a wealth of newly proposed techniques.

In this paper, we make the interested reader aware of the possibilities of feature selection, providing a basic taxonomy of feature selection techniques, and discussing their use, variety and potential in a number of both common as well as upcoming bioinformatics applications.

Companion website: `http://bioinformatics.psb.ugent.be/supplementary_data/yvsae/fsreview`

1 INTRODUCTION

During the last decade, the motivation for applying feature selection (FS) techniques in bioinformatics has shifted from being an illustrative example to becoming a real prerequisite for model building. In particular, the high dimensional nature of many modelling tasks in bioinformatics, going from sequence analysis over microarray analysis to spectral analyses and literature mining has given rise to a wealth of feature selection techniques being presented in the field.

In this review, we focus on the application of feature selection techniques. In contrast to other dimensionality reduction techniques like those based on projection (e.g. principal component analysis) or compression (e.g. using information theory), feature selection techniques do not alter the original representation of the variables, but merely select a subset of them. Thus, they preserve the original semantics of the variables, hence offering the advantage of interpretability by a domain expert.

While feature selection can be applied to both supervised and unsupervised learning, we focus here on the problem of supervised learning (classification), where the class labels are known beforehand. The interesting topic of feature selection for unsupervised learning (clustering) is a more complex issue, and research into this field is recently getting more attention in several communities [79, 122].

The main aim of this review is to make practitioners aware of the benefits, and in some cases even the necessity of applying feature selection techniques. Therefore, we provide an overview of the different feature selection techniques for classification: we illustrate them by reviewing the most important application fields in the bioinformatics domain, highlighting the efforts done by the bioinformatics community in developing novel and adapted procedures. Finally, we also point the interested reader to some useful data mining and bioinformatics software packages that can be used for feature selection.

2 FEATURE SELECTION TECHNIQUES

As many pattern recognition techniques were originally not designed to cope with large amounts of irrelevant features, combining them with FS techniques has become a necessity in many applications [43, 78, 79]. The objectives of feature selection are manifold, the most important ones being: a) to avoid overfitting and improve model performance, i.e. prediction performance in the case of supervised classification and better cluster detection in the case of clustering, b) to provide faster and more cost-effective models, and c) to gain a deeper insight into the underlying processes that generated the data. However, the advantages of feature selection techniques come at a certain price, as the search for a subset of relevant features introduces an additional layer of complexity in the modeling task. Instead of just optimizing the parameters of the model for the full feature subset, we now need to find the optimal model parameters for the optimal feature subset, as there is no guarantee that the optimal parameters for the full feature set are equally optimal for the optimal feature subset [20]. As a result, the search in the model hypothesis space is augmented by another dimension: the one of finding the optimal subset of relevant features. Feature selection techniques differ from each other in the way they incorporate this search in the added space of feature subsets in the model selection.

In the context of classification, feature selection techniques can be organized into three categories, depending on how they combine the feature selection search with the construction of the classification model: filter methods, wrapper methods, and embedded methods. Table 1 provides a common taxonomy of feature selection methods, showing for each technique the most prominent advantages and disadvantages, as well as some examples of the most influential techniques.

Filter techniques assess the relevance of features by looking only at the intrinsic properties of the data. In most cases a feature relevance score is calculated, and low scoring features are removed. Afterwards, this subset of features is presented as input to the classification algorithm. Advantages of filter techniques are that they easily scale to very high-dimensional datasets, they are computationally simple and fast, and they are independent of the classification algorithm. As a result, feature selection needs to be performed only once, and then different classifiers can be evaluated. A common disadvantage of filter methods is that they ignore the interaction with the classifier (the search in the feature subset space is separated from the search in the hypothesis space), and that most proposed techniques are univariate. This means that each feature is

© The Author (2007). Published by Oxford University Press. All rights reserved. For Permissions, please email: journals.permissions@oxfordjournals.org

Table 1. A taxonomy of feature selection techniques. For each feature selection type, we highlight a set of characteristics which can guide the choice for a technique suited to the goals and resources of practitioners in the field.

	Model search		Advantages	Disadvantages	Examples
Filter	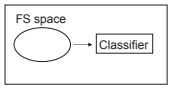	Univariate	Fast Scalable Independent of the classifier	Ignores feature dependencies Ignores interaction with the classifier	Chi-square Euclidean distance t-test Information gain, Gain ratio [6]
		Multivariate	Models feature dependencies Independent of the classifier Better computational complexity than wrapper methods	Slower than univariate techniques Less scalable than univariate techniques Ignores interaction with the classifier	Correlation based feature selection (CFS) [45] Markov blanket filter (MBF) [62] Fast correlation based feature selection (FCBF) [136]
Wrapper	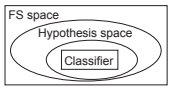	Deterministic	Simple Interacts with the classifier Models feature dependencies Less computationally intensive than randomized methods	Risk of over fitting More prone than randomized algorithms to getting stuck in a local optimum (greedy search) Classifier dependent selection	Sequential forward selection (SFS) [60] Sequential backward elimination (SBE) [60] Plus q take-away r [33] Beam search [106]
		Randomized	Less prone to local optima Interacts with the classifier Models feature dependencies	Computationally intensive Classifier dependent selection Higher risk of overfitting than deterministic algorithms	Simulated annealing Randomized hill climbing [110] Genetic algorithms [50] Estimation of distribution algorithms [52]
Embedded	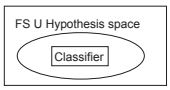		Interacts with the classifier Better computational complexity than wrapper methods Models feature dependencies	Classifier dependent selection	Decision trees Weighted naive Bayes [28] Feature selection using the weight vector of SVM [44, 125]

considered separately, thereby ignoring feature dependencies, which may lead to worse classification performance when compared to other types of feature selection techniques. In order to overcome the problem of ignoring feature dependencies, a number of multivariate filter techniques were introduced, aiming at the incorporation of feature dependencies to some degree.

Whereas filter techniques treat the problem of finding a good feature subset independently of the model selection step, *wrapper methods* embed the model hypothesis search within the feature subset search. In this setup, a search procedure in the space of possible feature subsets is defined, and various subsets of features are generated and evaluated. The evaluation of a specific subset of features is obtained by training and testing a specific classification model, rendering this approach tailored to a specific classification algorithm. To search the space of all feature subsets, a search algorithm is then "wrapped" around the classification model. However, as the space of feature subsets grows exponentially with the number of features, heuristic search methods are used to guide the search for an optimal subset. These search methods can be divided in two classes: deterministic and randomized search algorithms. Advantages of wrapper approaches include the interaction between feature subset search and model selection, and the ability to take into account feature dependencies. A common drawback of these techniques is that they have a higher risk of overfitting than filter techniques and are very computationally intensive, especially if building the classifier has a high computational cost.

In a third class of feature selection techniques, termed *embedded techniques*, the search for an optimal subset of features is built into the classifier construction, and can be seen as a search in the combined space of feature subsets and hypotheses. Just like wrapper approaches, embedded approaches are thus specific to a given learning algorithm. Embedded methods have the advantage that they include the interaction with the classification model, while at the same time being far less computationally intensive than wrapper methods.

3 APPLICATIONS IN BIOINFORMATICS

3.1 Feature selection for sequence analysis

Sequence analysis has a long standing tradition in bioinformatics. In the context of feature selection, two types of problems can be distinguished: content and signal analysis. Content analysis focuses on the broad characteristics of a sequence, such as tendency to code for proteins or fulfillment of a certain biological function. Signal analysis on the other hand focuses on the identification of important motifs in the sequence, such as gene structural elements or regulatory elements.

Apart from the basic features that just represent the nucleotide or amino acid at each position in a sequence, many other features, such as higher order combinations of these building blocks (e.g. k-mer patterns) can be derived, their number growing exponentially with the pattern length k. As many of them will be irrelevant or redundant, feature selection techniques are then applied to focus on the subset of relevant variables.

3.1.1 Content analysis The prediction of subsequences that code for proteins (coding potential prediction) has been a focus of interest since the early days of bioinformatics. Because many features can be extracted from a sequence, and most dependencies occur between adjacent positions, many variations of Markov models were developped. To deal with the high amount of possible features, and the often limited amount of samples, [101] introduced the interpolated Markov model (IMM), which used interpolation between different orders of the Markov model to deal with small sample sizes, and a filter method (Chi-square) to select only relevant features. In further work, [24] extended the IMM framework

to also deal with non-adjacent feature dependencies, resulting in the interpolated context model (ICM), which crosses a Bayesian decision tree with a filter method (Chi-square) to assess feature relevance. Recently, the avenue of FS techniques for coding potential prediction was further pursued by [100], who combined different measures of coding potential prediction, and then used the Markov blanket multivariate filter approach (MBF) to retain only the relevant ones.

A second class of techniques focuses on the prediction of protein function from sequence. The early work of [16], who combined a genetic algorithm in combination with the Gamma test to score feature subsets for classification of large subunits of rRNA, inspired researchers to use FS techniques to focus on important subsets of amino acids that relate to the protein's functional class [1]. An interesting technique is described in [137], using selective kernel scaling for support vector machines (SVM) as a way to asses feature weights, and subsequently remove features with low weights.

The use of FS techniques in the domain of sequence analysis is also emerging in a number of more recent applications, such as the recognition of promoter regions [18], and the prediction of microRNA targets [59].

3.1.2 Signal analysis Many sequence analysis methodologies involve the recognition of short, more or less conserved signals in the sequence, representing mainly binding sites for various proteins or protein complexes. A common approach to find regulatory motifs, is to relate motifs to gene expression levels using a regression approach. Feature selection can then be used to search for the motifs that maximize the fit to the regression model [58, 116]. In [109], a classification approach is chosen to find discriminative motifs. The method is inspired by [7] who use the threshold number of misclassification (TNoM, see further in the section on microarray analysis) to score genes for relevance to tissue classification. From the TNoM score, a p-value is calculated that represents the significance of each motif. Motifs are then sorted according to their p-value.

Another line of research is performed in the context of the gene prediction setting, where structural elements such as the translation initiation site (TIS) and splice sites are modelled as specific classification problems. The problem of feature selection for structural element recognition was pioneered in [23] for the problem of splice site prediction, combining a sequential backward method together with an embedded SVM evaluation criterion to assess feature relevance. In [99] an estimation of distribution algorithm (EDA, a generalization of genetic algorithms) was used to gain more insight in the relevant features for splice site prediction. Similarly, the prediction of TIS is a suitable problem to apply feature selection techniques. In [76], the authors demonstrate the advantages of using feature selection for this problem, using the feature-class entropy as a filter measure to remove irrelevant features.

In future research, FS techniques can be expected to be useful for a number of challenging prediction tasks, such as identifying relevant features related to alternative splice sites and alternative TIS.

3.2 Feature selection for microarray analysis

During the last decade, the advent of microarray datasets stimulated a new line of research in bioinformatics. Microarray data pose a great challenge for computational techniques, because of their large dimensionality (up to several tens of thousands of genes) and their small sample sizes [112]. Furthermore, additional experimental complications like noise and variability render the analysis of microarray data an exciting domain.

In order to deal with these particular characteristics of microarray data, the obvious need for dimension reduction techniques was realized [2, 7, 40, 97], and soon their application became a de facto standard in the field. Whereas in 2001, the field of microarray analysis was still claimed to be in its infancy [31], a considerable and valuable effort has since been done to contribute new and adapt known FS methodologies [53]. A general overview of the most influential techniques, organized according to the general FS taxonomy of Section 2, is shown in Table 2.

3.2.1 The univariate filter paradigm: simple yet efficient

Because of the high dimensionality of most microarray analyses, fast and efficient FS techniques such as univariate filter methods have attracted most attention. The prevalence of these univariate techniques has dominated the field, and up to now comparative evaluations of different classification and FS techniques over DNA microarray datasets only focused on the univariate case [29, 64, 72, 113]. This domination of the univariate approach can be explained by a number of reasons:

- the output provided by univariate feature rankings is intuitive and easy to understand;
- the gene ranking output could fulfill the objectives and expectations that bio-domain experts have when wanting to subsequently validate the result by laboratory techniques or in order to explore literature searches. The experts could not feel the need for selection techniques that take into account gene interactions;
- the possible unawareness of subgroups of gene expression domain experts about the existence of data analysis techniques to select genes in a multivariate way;
- the extra computation time needed by multivariate gene selection techniques.

Some of the simplest heuristics for the identification of differentially expressed genes include setting a threshold on the observed fold-change differences in gene expression between the states under study, and the detection of the threshold point in each gene that minimizes the number of training sample misclassification (threshold number of misclassification, TNoM [7]). However, a wide range of new or adapted univariate feature ranking techniques has since then been developped. These techniques can be divided into two classes: parametric and model-free methods (see Table 2).

Parametric methods assume a given distribution from which the samples (observations) have been generated. The two sample t-test and ANOVA are among the most widely used techniques in microarray studies, although the usage of their basic form, possibly without justification of their main assumptions, is not advisible [53]. Modifications of the standard t-test to better deal with the small sample size and inherent noise of gene expression datasets include a number of t- or t-test like statistics (differing primarily in the way the variance is estimated) and a number of Bayesian frameworks [4, 35]. Although Gaussian assumptions have dominated the field, other types of parametrical approaches can also be found in the literature, such as regression modelling approaches [117] and Gamma distribution models [85].

Due to the uncertainty about the true underlying distribution of

Table 2. Key references for each type of feature selection technique in the microarray domain.

Filter methods			Wrapper methods	Embedded methods
Univariate		Multivariate		
Parametric	Model-free			
t-test [53] ANOVA [53] Bayesian [4, 35] Regression [117] Gamma [85]	Wilcoxon rank sum [117] BSS/WSS [29] Rank products [12] Random permutations [31, 87, 88, 121] TNoM [7]	Bivariate [10] CFS [124, 131] MRMR [26] USC [132] Markov blanket [38, 82, 128]	Sequential search [51, 129] Genetic algorithms [56, 71, 86] Estimation of distribution algorithms [9]	Random forest [25, 55] Weight vector of SVM [44] Weights of logistic regression [81]

many gene expression scenarios, and the difficulties to validate distributional assumptions because of small sample sizes, non-parametric or model-free methods have been widely proposed as an attractive alternative to make less stringent distributional assumptions [120]. Many model-free metrics, frequently borrowed from the statistics field, have demonstrated their usefulness in many gene expression studies, including the Wilcoxon rank-sum test [117], the between-within classes sum of squares (BSS/WSS) [29] and the rank products method [12].

A specific class of model-free methods estimates the reference distribution of the statistic using random permutations of the data, allowing the computation of a model-free version of the associated parametric tests. These techniques have emerged as a solid alternative to deal with the specificities of DNA microarray data, and do not depend on strong parametric assumptions [31, 87, 88, 121]. Their permutation principle partly alleviates the problem of small sample sizes in microarray studies, enhancing the robustness against outliers.

We also mention promising types of non-parametric metrics which, instead of trying to identify differentially expressed genes at the whole population level (e.g. comparison of sample means), are able to capture genes which are significantly disregulated in only a subset of samples [80, 89]. These types of methods offer a more patient specific approach for the identification of markers, and can select genes exhibiting complex patterns that are missed by metrics that work under the classical comparison of two prelabeled phenotypic groups. In addition, we also point out the importance of procedures for controlling the different types of errors that arise in this complex multiple testing scenario of thousands of genes [30, 92, 93, 114], with a special focus on contributions for controlling the false discovery rate (FDR).

3.2.2 Towards more advanced models: the multivariate paradigm for filter, wrapper and embedded techniques

Univariate selection methods have certain restrictions and may lead to less accurate classifiers by, for example, not taking into account gene-gene interactions. Thus, researchers have proposed techniques that try to capture these correlations between genes.

The application of multivariate filter methods ranges from simple bivariate interactions [10] towards more advanced solutions exploring higher order interactions, such as correlation based feature selection (CFS) [124, 131] and several variants of the Markov blanket filter method [38, 82, 128]. The Minimum Redundancy - Maximum Relevance (MRMR) [26] and Uncorrelated Shrunken Centroid (USC) [132] algorithms are two other solid multivariate filter procedures, highlighting the advantage of using multivariate methods over univariate procedures in the gene expression domain.

Feature selection using wrapper or embedded methods offers an alternative way to perform a multivariate gene subset selection, incoporating the classifier's bias into the search and thus offering an opportunity to construct more accurate classifiers. In the context of microarray analysis, most wrapper methods use population based, randomized search heuristics [9, 56, 71, 86], although also a few examples use sequential search techniques [51, 129]. An interesting hybrid filter-wrapper approach is introduced in [98], crossing a univariately pre-ordered gene ranking with an incrementally augmenting wrapper method.

Another characteristic of any wrapper procedure concerns the scoring function used to evaluate each gene subset found. As the 0-1 accuracy measure allows for comparison with previous works, the vast majority of papers uses this measure. However, recent proposals advocate the use of methods for the approximation of the area under the ROC curve [81], or the optimization of the LASSO (Least Absolute Shrinkage and Selection Operator) model [39]. ROC curves certainly provide an interesting evaluation measure, especially suited to the demand for screening different types of errors in many biomedical scenarios.

The embedded capacity of several classifiers to discard input features and thus propose a subset of discriminative genes, has been exploited by several authors. Examples include the use of random forests (a classifier that combines many single decision trees) in an embedded way to calculate the importance of each gene [25, 55]. Another line of embedded FS techniques uses the weights of each feature in linear classifiers such as SVMs [44] and logistic regression [81]. These weights are used to reflect the relevance of each gene in a multivariate way, and thus allow for the removal of genes with very small weights.

Partially due to the higher computational complexity of wrapper and to a lesser degree embedded approaches, these techniques have not received as much interest as filter proposals. However, an advisable practice is to pre-reduce the search space using a univariate filter method, and only then apply wrapper or embedded methods, hence fitting the computation time to the available resources.

3.3 Mass spectra analysis

Mass spectrometry technology (MS) is emerging as a new and attractive framework for disease diagnosis and protein-based biomarker profiling [91]. A mass spectrum sample is characterized by thousands of different mass/charge (m/z) ratios on the x-axis, each with their corresponding signal intensity value on the y-axis. A typical MALDI-TOF low-resolution proteomic profile can contain

Table 3. Key references for each type of feature selection technique in the domain of mass spectrometry.

	Univariate		Multivariate
	Parametric	Model-free	
Filter	t-test [77, 127] F-test [8]	Peak Probability Contrast [118] Kolmogorov-Smirnov test [135]	CFS [77] Relief-F [94]
Wrapper	Genetic algorithms [70, 90] Nature inspired [95, 96]		
Embedded	Random forest/decision tree [37, 127] Weight vector of SVM [57, 138, 94] Neural network [5]		

up to $15,500$ data points in the spectrum between 500 and $20,000$ m/z, and the number of points even grows using higher resolution instruments.

For data mining and bioinformatics purposes, it can initially be assumed that each m/z ratio represents a distinct variable whose value is the intensity. As Somorjai *et al.* [112] explain, the data analysis step is severely constrained by both high dimensional input spaces and their inherent sparseness, just as it is the case with gene expression datasets. Although the amount of publications on mass spectrometry based data mining is not comparable to the level of maturity reached in the microarray analysis domain, an interesting collection of methods has been presented in the last 4-5 years (see [49, 105] for recent reviews) since the pioneering work of Petricoin *et al.* [90].

Starting from the raw data, and after an inital step to reduce noise and normalize the spectra from different samples [19], the following crucial step is to extract the variables that will constitute the initial pool of candidate discriminative features. Some studies employ the simplest approach of considering every measured value as a predictive feature, thus applying FS techniques over initial huge pools of about $15,000$ variables [70, 90], up to around $100,000$ variables [5]. On the other hand, a great deal of the current studies performs aggressive feature extraction procedures using elaborated peak detection and alignment techniques (see [19, 49, 105] for a detailed description of these techniques). These procedures tend to seed the dimensionality from which supervised FS techniques will start their work in less than 500 variables [8, 96, 118]. A feature extraction step is thus advisable to set the computational costs of many FS techniques to a feasible size in these MS scenarios. Table 3 presents an overview of FS techniques used in the domain of mass spectrometry. Similar to the domain of microarray analysis, univariate filter techniques seem to be the most common techniques used, although the use of embedded techniques is certainly emerging as an alternative. Although the t-test maintains a high level of popularity [77, 127], other parametric measures (such as F-test [8]), and a notable variety of non-parametric scores [118, 135] have also been used in several MS studies. Multivariate filter techniques on the other hand, are still somewhat underrepresented [77, 94].

Wrapper approaches have demonstrated their usefulness in MS studies by a group of influential works. Different types of population based randomized heuristics are used as search engines in the major part of these papers: genetic algorithms [70, 90], particle swarm optimization [95] and ant colony procedures [96]. It is worth noting that while the first two references start the search procedure in $\approx 15,000$ dimensions by considering each m/z ratio as an initial predictive feature, aggressive peak detection and alignment processes reduce the initial dimension to about 300 variables in the last two references [95, 96].

An increasing number of papers uses the embedded capacity of several classifiers to discard input features. Variations of the popular method originally proposed for gene expression domains by Guyon *et al.* [44], using the weights of the variables in the SVM-formulation to discard features with small weights, have been broadly and successfully applied in the MS domain [57, 94, 138]. Based on a similar framework, the weights of the input masses in a neural network classifier have been used to rank the features' importance in Ball *et al.* [5]. The embedded capacity of random forests [127] and other types of decision tree based algorithms [37] constitutes an alternative embedded FS strategy.

4 DEALING WITH SMALL SAMPLE DOMAINS

Small sample sizes, and their inherent risk of imprecision and overfitting, pose a great challenge for many modelling problems in bioinformatics [11, 84, 108]. In the context of feature selection, two initiatives have emerged in response to this novel experimental situation: the use of adequate evaluation criteria, and the use of stable and robust feature selection models.

4.1 Adequate evaluation criteria

Several papers have warned about the substantial number of applications not performing an independent and honest validation of the reported accuracy percentages [3, 113, 112]. In such cases, authors often select a discriminative subset of features using the whole dataset. The accuracy of the final classification model is estimated using this subset, thus testing the discrimination rule on samples that were already used to propose the final subset of features. We feel that the need for an external feature selection process in training the classification rule at each stage of the accuracy estimation procedure is gaining space in the bioinformatics community practices.

Furthermore, novel predictive accuracy estimation methods with promising characteristics, such as bolstered error estimation [107], have emerged to deal with the specificities of small sample domains.

4.2 Ensemble feature selection approaches

Instead of choosing one particular FS method, and accepting its outcome as the final subset, different FS methods can be combined using *ensemble FS approaches*. Based on the evidence that there is often not a single universally optimal feature selection technique [130], and due to the possible existence of more than one subset of features that discriminates the data equally well [133], model combination approaches such as boosting have been adapted to improve the robustness and stability of final, discriminative methods [7, 29].

Novel ensemble techniques in the microarray and mass spectrometry domains include averaging over multiple single feature subsets [69, 73], integrating a collection of univariate differential gene expression purpose statistics via a distance synthesis scheme [130], using different runs of a genetic algorithm to asses relative

importancies of each feature [70, 71], computing the Kolmogorov-Smirnov test in different bootstrap samples to assign a probability of being selected to each peak [134], and a number of Bayesian averaging approaches [65, 133]. Furthermore, methods based on a collection of decision trees (e.g. random forests) can be used in an ensemble FS way to assess the relevance of each feature [25, 37, 55, 127].

Although the use of ensemble approaches requires additional computational resources, we would like to point out that they offer an advisable framework to deal with small sample domains, provided the extra computational resources are affordable.

5 FEATURE SELECTION IN UPCOMING DOMAINS

5.1 Single nucleotide polymorphism analysis

Single nucleotide polymorphisms (SNPs) are mutations at a single nucleotide position that occurred during evolution and were passed on through heredity, accounting for most of the genetic variation among different individuals. SNPs are at the forefront of many disease-gene association studies, their number being estimated at about 7 million in the human genome [63]. Thus, selecting a subset of SNPs that is sufficiently informative but still small enough to reduce the genotyping overhead is an important step towards disease-gene association. Typically, the number of SNPs considered is not higher than tens of thousands with sample sizes of about one hundred.

Several computational methods for htSNP selection (haplotype SNPs; a set of SNPs located on one chromosome) have been proposed in the past few years. One approach is based on the hypothesis that the human genome can be viewed as a set of discrete blocks that only share a very small set of common haplotypes [21]. This approach aims to identify a subset of SNPs that can either distinguish all the common haplotypes [36], or at least explain a certain percentage of them. Another common htSNP selection approach is based on pairwise associations of SNPs, and tries to select a set of htSNPs such that each of the SNPs on a haplotype is highly associated with one of the htSNPs [15]. A third approach considers htSNPs as a subset of all SNPs, from which the remaining SNPs can be reconstructed [46, 66, 75]. The idea is to select htSNPs based on how well they predict the remaining set of the unselected SNPs.

When the haplotype structure in the target region is unknown, a widely used approach is to choose markers at regular intervals [67], given either the number of SNPs to choose or the desired interval. In [74] an ensemble approach is successfully applied to the identification of relevant SNPs for alcoholism, while [41] propose a robust feature selection technique based on a hybrid between a genetic algorithm and an SVM. The Relief-F feature selection algorithm, in conjunction with three classification algorithms (k-NN, SVM and naive Bayes) has been proposed in [123]. Genetic algorithms have been applied to the search of the best subset of SNPs, evaluating them with a multivariate filter (CFS), and also in a wrapper manner (with a decision tree as supervised classification paradigm) [103]. The multiple linear regression SNP prediction algorithm [48] predicts a complete genotype based on the values of its informative SNPs (selected with a stepwise tag selection algorithm), their positions among all SNPS, and a sample of complete genotypes. In [104] the tag SNP selection method allows to specify variable tagging thresholds, based on correlations, for different SNPs.

5.2 Text and literature mining

Text and literature mining is emerging as a promising area for data mining in biology [17, 54]. One important representation of text and documents is the so-called bag-of-words (BOW) representation, where each word in the text represents one variable, and its value consists of the frequency of the specific word in the text. It goes without saying that such a representation of the text may lead to very high dimensional datasets, pointing out the need for feature selection techniques.

Although the application of feature selection techniques is common in the field of text classification (see e.g. [34] for a review), the application in the biomedical domain is still in its infancy. Some examples of FS techniques in the biomedical domain include the work of Dobrokhotov *et al.* [27], who use the Kullback-Leibler divergence as a univariate filter method to find discriminating words in a medical annotation task, the work of Eom and Zhang [32] who use symmetrical uncertainty (an entropy based filter method) for identifying relevant features for protein interaction discovery, and the work of Han *et al.* [47], which discusses the use of feature selection for a document classification task.

It can be expected that, for tasks such as biomedical document clustering and classification, the large number of feature selection techniques that were already developed in the text mining community will be of practical use for researchers in biomedical literature mining [17].

6 FS SOFTWARE PACKAGES

In order to provide the interested reader with some pointers to existing software packages, Table 4 shows an overview of existing software implementing a variety of feature selection methods. All software packages mentioned are free for academic use, and the software is organized into four sections: general purpose FS techniques, techniques tailored to the domain of microarray analysis, techniques specific to the domain of mass spectra analysis, and techniques to handle SNP selection. For each software package, the main reference, implementation language and website is shown. In addition to these publicly available packages, we also provide a companion website of this work (see the Abstract for the location). On this website, the publications are indexed according to the FS technique used, a number of keywords accompanying each reference to understand its FS methodological contributions.

7 CONCLUSIONS AND FUTURE PERSPECTIVES

In this paper, we reviewed the main contributions of feature selection research in a set of well-known bioinformatics applications. Two main issues emerge as common problems in the bioinformatics domain: the large input dimensionality, and the small sample sizes. To deal with these problems, a wealth of FS techniques has been designed by researchers in bioinformatics, machine learning and data mining.

A large and fruitful effort has been performed during the last years in the adaptation and proposal of univariate filter FS techniques. In general, we observe that many researchers in the field still think that filter FS approaches are only restricted to univariate approaches. The proposal of multivariate selection algorithms can be considered as

Table 4. Software for feature selection.

General purpose FS software			
WEKA	Java	[126]	http://www.cs.waikato.ac.nz/ml/weka
Fast Correlation Based Filter	Java	[136]	http://www.public.asu.edu/~huanliu/FCBF/FCBFsoftware.html
Feature Selection Book	Ansi C	[78]	http://www.public.asu.edu/~huanliu/Fsbook
MLC++	C++	[61]	http://www.sgi.com/tech/mlc
Spider	Matlab	-	http://www.kyb.tuebingen.mpg.de/bs/people/spider
SVM and Kernel Methods Matlab Toolbox	Matlab	[14]	http://asi.insa-rouen.fr/~arakotom/toolbox/index
Microarray analysis FS software			
SAM	R, Excel	[121]	http://www-stat.stanford.edu/~tibs/SAM/
GALGO	R	[119]	http://www.bip.bham.ac.uk/bioinf/galgo.html
PCP	C, C++	[13]	http://pcp.sourceforge.net
GA-KNN	C	[71]	http://dir.niehs.nih.gov/microarray/datamining/
Rankgene	C	[115]	http://genomics10.bu.edu/yangsu/rankgene/
EDGE	R	[68]	http://www.biostat.washington.edu/software/jstorey/edge/
GEPAS-Prophet	Perl, C	[83]	http://prophet.bioinfo.cipf.es/
DEDS (Bioconductor)	R	[130]	http://www.bioconductor.org/
RankProd (Bioconductor)	R	[12]	http://www.bioconductor.org/
Limma (Bioconductor)	R	[111]	http://www.bioconductor.org/
Multtest (Bioconductor)	R	[30]	http://www.bioconductor.org/
Nudge (Bioconductor)	R	[22]	http://www.bioconductor.org/
Qvalue (Bioconductor)	R	[114]	http://www.bioconductor.org/
twilight (Bioconductor)	R	[102]	http://www.bioconductor.org/
ComparativeMarkerSelection (GenePattern)	JAVA, R	[42]	http://www.broad.mit.edu/genepattern
Mass spectra analysis FS software			
GA-KNN	C	[70]	http://dir.niehs.nih.gov/microarray/datamining/
R-SVM	R, C, C++	[138]	http://www.hsph.harvard.edu/bioinfocore/RSVMhome/R-SVM.html
SNP analysis FS software			
CHOISS	C++, Perl	[67]	http://biochem.kaist.ac.kr/choiss.htm
MLR-tagging	C	[48]	http://alla.cs.gsu.ed/~software/tagging/tagging.html
WCLUSTAG	JAVA	[104]	http://bioinfo.hku.hk/wclustag

one of the most promising future lines of work for the bioinformatics community.

A second line of future research is the development of especially fitted ensemble FS approaches to enhance the robustness of the finally selected feature subsets. We feel that, in order to alleviate the actual small sample sizes of the majority of bioinformatics applications, the further development of such techniques, combined with appropriate evaluation criteria, constitutes an interesting direction for future FS research.

Other interesting opportunities for future FS research will be the extension towards upcoming bioinformatics domains such as SNPs, text and literature mining, and the combination of heterogeneous data sources. While in these domains, the FS component is not yet as central as e.g. in gene expression or MS areas, we believe that its application will become essential in dealing with the high dimensional character of these applications.

To conclude, we would like to note that, in order to maintain an appropriate size of the paper, we had to limit the number of referenced studies. We therefore apologize to the authors of papers that were not cited in this work.

8 ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers for their constructive comments, which significantly improved the quality of this review. This work was supported by BOF grant 01P10306 from Ghent University to Y.S., and the SAIOTEK and ETORTEK programs of the Basque Government and project TIN2005-03824 of the Spanish Ministry of Education and Science to I.I. and P.L.

REFERENCES

[1]A. Al-Shahib, R. Breitling, and D. Gilbert. Feature selection and the class imbalance problem in predicting protein function from sequence. *Applied Bioinformatics*, 4(3):195–203, 2005.

[2]U. Alon, N. Barkai, K. G. D. Notterman, S. Ybarra, D. Mack, and A. Levine. Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissues probed by oligonucleotide arrays. In *Proceedings of the National Academy of Sciences, USA*, volume 96, pages 6745–6750, 1999.

[3]C. Ambroise and G. McLachlan. Selection bias in gene extraction on the basis of microarray gene-expression data. In *Proceedings of the National Academy of Sciences*, volume 99, pages 6562–6566, 2002.

[4]P. Baldi and A. Long. A Bayesian framework for the analysis of microarray expression data: regularized t-test and statistical inferences of gene changes. *Bioinformatics*, 17(6):509–516, 2001.

[5]G. Ball, S. Mian, F. Holding, R. Allibone, J. Lowe, S. Ali, G. Li, S. McCardle, I. Ellis, C. Creaser, and R. Rees. An integrated approach utilizing artificial neural networks and SELDI mass spectrometry for the classification of human tumours and rapid identification of potential biomarkers. *Bioinformatics*, 18(3):395–404, 2002.

[6]M. Ben-Bassat. Pattern recognition and reduction of dimensionality. In P. Krishnaiah and L. Kanal, editors, *Handbook of Statistics II*, volume 1, pages 773–791. North-Holland, 1982. Amsterdam.

[7]A. Ben-Dor, L. Bruhn, N. Friedman, I. Nachman, M. Schummer, and Z. Yakhini. Tissue classification with gene expression profiles. *Journal of Computational Biology*, 7(3-4):559–584, 2000.

[8]G. Bhanot, G. Alexe, B. Venkataraghavan, and A. Levine. A robust meta-classification strategy for cancer detection from MS data. *Proteomics*, 6(2):592–604, 2006.

[9]R. Blanco, P. Larrañaga, I. Inza, and B. Sierra. Gene selection for cancer classification using wrapper approaches. *International Journal of Pattern Recognition and Artificial Intelligence*, 18(8), 2004.

[10]T. Bø and I. Jonassen. New feature subset selection procedures for classification of expression profiles. *Genome Biology*, 3(4):research0017.1–research0017.11, 2002.

[11]U. Braga-Neto and E. Dougherty. Is cross-validation valid for small-sample microarray classification? *Bioinformatics*, 20(3):374–380, 2004.

[12]R. Breitling, P. Armengaud, A. Amtmann, and P. Herzyk. Rank products: a simple, yet powerful, new method to detect differentially regulated genes in replicated microarray experiments. *FEBS Letters*, 573:83–92, 2004.

[13]L. Buturovic. PCP: a program for supervised classification of gene expression profiles. *Bioinformatics*, 22(2):245–247, 2005.

[14]S. Canu, Y. Grandvalet, and A. Rakotomamonjy. SVM and Kernel Methods Matlab Toolbox. In *Perception Systèmes et Information, INSA de Rouen, Rouen, France*, 2003.

[15]C. Carlson, M. Eberle, M. Rieder, Q. Yi, L. Kruglyak, and D. Nickerson. Selecting a maximally informative set of single-nucleotide polymorphisms for association analyses using linkage disequilibrium. *American Journal of Human Genetics*, 74:106–120, 2004.

[16]N. Chuzhanova, A. Jones, and S. Margetts. Feature selection for genetic sequence classification. *Bioinformatics*, 14(2):139–143, 1998.

[17]A. Cohen and W. Hersch. A survey of current work in biomedical text mining. *Briefings in Bioinformatics*, 6(1):57–71, 2005.

[18]P. Conilione and D. Wang. A comparative study on feature selection for E.coli promoter recognition. *International Journal of Information Technology*, 11:54–66, 2005.

[19]K. Coombes, K. Baggerly, and J. Morris. Pre-processing mass spectometry data. In M. Dubitzky, M. Granzow, and D. Berrar, editors, *Fundamentals of Data Mining in Genomics and Proteomics*, pages 79–99. Kluwer, 2007.

[20]W. Daelemans, V. Hoste, F. De Meulder, and B. Naudts. Combined optimization of feature selection and algorithm parameter interaction in machine learning of language. In *Proceedings of the 14th European Conference on Machine Learning (ECML-2003)*, pages 84–95, 2003.

[21]M. Daly, J. Rioux, S. Schaffner, T. Hudson, and E. Lander. High-resolution haplotype structure in the human genome. *Nature Genetics*, 29:229–232, 2001.

[22]N. Dean and A. Raftery. Normal uniform mixture differential gene expression detection in cDNA microarrays. *BMC Bioinformatics*, 6(173), 2005.

[23]S. Degroeve, B. De Baets, Y. Van de Peer, and P. Rouzé. Feature subset selection for splice site prediction. *Bioinformatics*, 18 Supp.2:75–83, 2002.

[24]A. Delcher, D. Harnon, S. Kasif, O. White, and S. Salzberg. Improved microbial gene identification with GLIMMER. *Nucleic Acids Research*, 27:4636–4641, 1999.

[25]R. Díaz-Uriarte and S. Alvarez de Andrés. Gene selection and classification of microarray data using random forest. *BMC Bioinformatics*, 7(3), 2006.

[26]C. Ding and H. Peng. Minimum redundancy feature selection from microarray gene expression data. In *Proceedings of the IEEE Conference on Computational Systems Bioinformatics*, pages 523–528, 2003.

[27]P. Dobrokhotov, C. Goutte, A. Veuthey, and E. Gaussier. Combining NLP and probabilistic categorisation for document and term selection for Swiss-Prot medical annotation. *Bioinformatics*, 19 Supp.1:91–94, 2003.

[28]P. Duda, P. Hart, and D. Stork. *Pattern Classification*. Wiley, New York, 2001.

[29]S. Dudoit, J. Fridlyand, and T. Speed. Comparison of discriminant methods for the classification of tumors using gene expression data. *Journal of the American Statistical Association*, 97(457):77–87, 2002.

[30]S. Dudoit, J. Shaffer, and J. Boldrick. Multiple hypothesis testing in microarray experiments. *Statistical Science*, 18:7–103, 2003.

[31]B. Efron, R. Tibshirani, J. Storey, and V. Tusher. Empirical Bayes analysis of a microarray experiment. *Journal of the American Statistical Association*, 96(456):1151–1160, 2001.

[32]J. Eom and B. Zhang. PubMiner:machine learning-based text mining for biomedical information analysis. In *Lecture Notes in Artificial Intelligence*, volume 3192, pages 216–225, 2000.

[33]F. Ferri, P. Pudil, M. Hatef, and J. Kittler. *Pattern Recognition in Practice IV, Multiple Paradigms, Comparative Studies and Hybrid Systems*, chapter Comparative study of techniques for large-scale feature selection, pages 403–413. Elsevier, 1994.

[34]G. Forman. An extensive empirical study of feature selection metrics for text classification. *Journal of Machine Learning Research*, 3:1289–1305, 2003.

[35]R. Fox and M. Dimmic. A two-sample Bayesian t-test for microarray data. *BMC Bioinformatics*, 7(1):126, 2006.

[36]S. Gabriel, S. Schaffner, H. Nguyen, J. Moore, J. Roy, B. Blumenstiel, J. Higgins, M. DeFelice, A. Lochner, M. Faggart, S. Liu-Cordero, C. Rotimi, A. Adeyemo, R. Cooper, R. Ward, E. Lander, M. Daly, and D. Altshuler. The structure of haplotype blocks in the human genome. *Science*, 296:2225–2229, 2002.

[37]P. Geurts, M. Fillet, D. de Seny, M.-A. Meuwis, M. Malaise, M.-P. Merville, and L. Wehenkel. Proteomic mass spectra classification using decision tree based ensemble methods. *Bioinformatics*, 21(15):3138–3145, 2005.

[38]O. Gevaert, F. De Smet, D. Timmerman, Y. Moreau, and B. De Moor. Predicting the prognosis of breast cancer by integrating clinical and microarray data with Bayesian networks. *Bioinformatics*, 22(14):e184–e190, 2006.

[39]D. Ghosh and M. Chinnaiyan. Classification and selection of biomarkers in genomic data using LASSO. *Journal of Biomedicine and Biotechnology*, 2005(2):147–154, 2005.

[40]T. Golub, D. Slonim, P. Tamayo, C. Huard, M. Gaasenbeek, J. Mesirov, H. Coller, M. Loh, J. Downing, M. Caliguri, C. Bloomfield, and E. Lander. Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *Science*, 286:531–537, 1999.

[41]B. Gong, Z. Guo, J. Li, G. Zhu, S. Lv, S. Rao, and X. Li. Application of genetic algorithm – support vector machine hybrid for prediction of clinical phenotypes based on geneome-wide SNP profiles of sib pairs. In *Lecture Notes in Computer Science 3614*, pages 830–835. Springer, 2005.

[42]J. Gould, G. Getz, S. Monti, M. Reich, and J. Mesirov. Comparative gene marker selection suite. *Bioinformatics*, 22(15):1924–1925, 2006.

[43]I. Guyon and A. Elisseeff. An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003.

[44]I. Guyon, J. Weston, S. Barnhill, and V. Vapnik. Gene selection for cancer classification using support vector machines. *Machine Learning*, 46(1-3):389–422, 2002.

[45]M. Hall. *Correlation-based Feature Selection for Machine Learning*. PhD thesis, Department of Computer Science, Waikato University, New Zealand, 1999.

[46]E. Halperin, G. Kimmel, and R. Shamir. Tag SNP selection in genotype data for maximizing SNP prediction accuracy. *Bioinformatics*, 21(suppl. 1):i195–203, 2005.

[47]B. Han, Z. Obradovic, Z. Hu, C. Wu, and S. Vucetic. Substring selection for biomedical document classification. *Bioinformatics*, 22(17):2136–2142, 2006.

[48]J. He and A. Zelikovsky. MLR-tagging: informative SNP selection for unphased genotypes based on multiple linear regression. *Bioinformatics*, 22(20):2558–2561, 2006.

[49]M. Hilario, A. Kalousis, C. Pellegrini, and M. Muller. Processing and classification of protein mass spectra. *Mass Spectometry Reviews*, 25(3):409–449, 2006.

[50]J. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, 1975.

[51]I. Inza, P. Larrañaga, R. Blanco, and A. Cerrolaza. Filter versus wrapper gene selection approaches in DNA microarray domains. *Artificial Intelligence in Medicine*, 31(2):91–103, 2004.

[52]I. Inza, P. Larrañaga, R. Etxebarria, and B. Sierra. Feature subset selection by Bayesian networks based optimization. *Artifical Intelligence*, 123(1-2):157–184, 2000.

[53]P. Jafari and F. Azuaje. An assessment of recently published gene expression data analyses: reporting experimental design and statistical factors. *BMC Medical Informatics and Decision Making*, 6(1):27, 2006.

[54]L. Jensen, J. Saric, and P. Bork. Literature mining for the biologist: from information retrieval to biological discovery. *Nature Reviews Genetics*, 7(2):119–129, 2006.

[55]H. Jiang, Y. Deng, H.-S. Cheng, L. Tao, Q. Sha, J. Chen, C.-J. Tsai, and S. Zhang. Joint analysis of two microarray gene-expression data sets to select lung adenocarcinoma marker genes. *BMC Bioinformatics*, 5(81), 2004.

[56]T. Jirapech-Umpai and S. Aitken. Feature selection and classification for microarray data analysis: evolutionary methods for identifying predictive genes. *BMC Bioinformatics*, 6(148), 2005.

[57]K. Jong, E. Marchiori, M. Sebag, and A. van der Vaart. Feature selection in proteomic pattern data with support vector machines. In *Proceedings of the IEEE Symposium on Computational Intelligence in Bioinformatics and Computational*

Biology, pages 41–48, 2004.

[58]S. Keles, M. van der Laan, and M. Eisen. Identification of regulatory elements using a feature selection method. *Bioinformatics*, 18(9):1167–1175, 2002.

[59]S. Kim, J. Nam, J. Rhee, W. Lee, and B. Zhang. miTarget: microRNA target gene prediction using a support vector machine. *BMC Bioinformatics*, 7(411), 2006.

[60]J. Kittler. *Pattern Recognition and Signal Processing*, chapter Feature set search algorithms, pages 41–60. Sijthoff and Noordhoff, Alphen aan den Rijn, Netherlands, 1978.

[61]R. Kohavi, D. Sommerfield, and J. Dougherty. Data mining using MLC++: a machine learning library in C++. In *Tools with Artificial Intelligence*, pages 234–245. IEEE Computer Society Press, 1996.

[62]D. Koller and M. Sahami. Toward optimal feature selection. In *Proceedings of the Thirteenth International Conference on Machine Learning*, pages 284–292, Bari, Italy, 1996.

[63]L. Kruglyak and D. A. Nickerson. Variation in the spice of life. *Nature Genetics*, 27:234–236, 2001.

[64]J. Lee, J. Loo, M. Park, and S. Song. An extensive comparison of recent classification tools applied to microarray data. *Computational Statistics and Data Analysis*, 48:869–885, 2005.

[65]K. Lee, N. Sha, E. Dougherty, M. Vannucci, and B. Mallick. Gene selection: a Bayesian variable selection approach. *Bioinformatics*, 19(1):90–97, 2003.

[66]P. H. Lee and H. Shatkay. BNTagger: improved tagging SNP selection using Bayesian networks. *Bioinformatics*, 22(14):e211–e219, 2006.

[67]S. Lee and C. Kang. CHOISS for selection on single nucleotide polymorphism markers on interval regularity. *Bioinformatics*, 20(4):581–582, 2004.

[68]J. Leek, E. Monsen, A. Dabney, and J. Storey. EDGE: extraction and analysis of differential gene expression. *Bioinformatics*, 22(4):507–508, 2006.

[69]I. Levner. Feature selection and nearest centroid classification for protein mass spectrometry. *BMC Bioinformatics*, 6(68), 2005.

[70]L. Li, D. Umbach, P. Terry, and J. Taylor. Applications of the GA/KNN method to SELDI proteomics data. *Bioinformatics*, 20(10):1638–1640, 2004.

[71]L. Li, C. Weinberg, T. Darden, and L. Pedersen. Gene selection for sample classification based on gene expression data: study of sensitivity to choice of parameters of the GA/KNN method. *Bioinformatics*, 17(12):1131–1142, 2001.

[72]T. Li, C. Zhang, and M. Ogihara. A comparative study of feature selection and multiclass classification methods for tissue classification based on gene expression. *Bioinformatics*, 20(15):2429–2437, 2004.

[73]W. Li and Y. Yang. How many genes are needed for a discriminant microarray data analysis? In S. M. Lin and K. F. Johnson, editors, *Methods of Microarray Data Analysis. First Conference on Critical Assessment of Microarray Data Analysis, CAMDA2000*, pages 137–150, 2002.

[74]X. Li, S. Rao, W.Zhang, G. Zheng, W. Jiang, and L. Du. Large-scale ensemble decision analysis of sib-pair ibd profiles for identification of the relevant molecular signatures for alcoholism. In *Lecture Notes in Computer Science 3614*, pages 1184–1189. Springer, 2005.

[75]Z. Lin and R. B. Altman. Finding haplotype tagging SNPs by use of principal components analysis. *American Journal of Human Genetics*, 73:850–861, 2004.

[76]H. Liu, H. Han, J. Li, and L. Wong. Using amino acid patterns to accurately predict translation initiation sites. *In Silico Biology*, 4(3):255–269, 2004.

[77]H. Liu, J. Li, and L. Wong. A comparative study on feature selection and classification methods using gene expression profiles and proteomic patterns. *Genome Informatics*, 13:51–60, 2002.

[78]H. Liu and H. Motoda. *Feature Selection for Knowledge Discovery and Data Mining*. Kluwer Academic Publishers, 1998.

[79]H. Liu and L. Yu. Toward integrating feature selection algorithms for classification and clustering. *IEEE Transactions on Knowledge and Data Engineering*, 17(4):491–502, 2005.

[80]J. Lyons-Weiler, S. Patel, M. Becich, and T. Godfrey. Tests for finding complex patterns of differential expression in cancers: towards individualized medicine. *BMC Bioinformatics*, 5(110), 2004.

[81]S. Ma and J. Huang. Regularized ROC method for disease classification and biomarker selection with microarray data. *Bioinformatics*, 21(24):4356–4362, 2005.

[82]H. Mamitsuka. Selecting features in microarray classification using ROC curves. *Pattern Recognition*, 39:2393–2404, 2006.

[83]I. Medina, D. Montaner, J. Tárraga, and J. Dopazo. Prophet, a web-based tool for class prediction using microarray data. *Bioinformatics*, 23(3):390–391, 2007.

[84]A. Molinaro, R. Simon, and R. Pfeiffer. Prediction error estimation: a comparison of resampling methods. *Bioinformatics*, 21(15):3301–3307, 2005.

[85]M. Newton, C. Kendziorski, C. Richmond, C. Blattner, and K. Tsui. On differential variability of expression ratios: improving statistical inference about gene expression changes from microarray data. *Journal of Computational Biology*, 8:37–52, 2001.

[86]C. Ooi and P. Tan. Genetic algorithms applied to multi-class prediction for the analysis of gene expression data. *Bioinformatics*, 19(1):37–44, 2003.

[87]W. Pan. On the use of permutation in and the performance of a class of nonparametric methods to detect differential gene expression. *Bioinformatics*, 19(11):1333–1340, 2003.

[88]P. Park, M. Pagano, and M. Bonetti. A nonparametric scoring algorithm for identifying informative genes from microarray data. *Pacific Symposium on Biocomputing*, 6:52–63, 2001.

[89]P. Pavlidis and P. Poirazi. Individualized markers optimize class prediction of microarray data. *BMC Bioinformatics*, 7(1):345, 2006.

[90]E. Petricoin, A. Ardekani, B. Hitt, P. Levine, V. Fusaro, S. Steinberg, G. Mills, C. Simone, D. Fishman, E. Kohn, and L. Liotta. Use of proteomics patterns in serum to identify ovarian cancer. *The Lancet*, 359(9306):572–577, 2002.

[91]E. Petricoin and L. Liotta. Mass spectometry-based diagnostic: the upcoming revolution in disease detection. *Clinical Chemistry*, 49(4):533–534, 2003.

[92]A. Ploner, S. Calza, A. Gusnanto, and Y. Pawitan. Multidimensional local false discovery rate for microarray studies. *Bioinformatics*, 22(5):556–565, 2006.

[93]S. Pounds and C. Cheng. Improving false discovery rate estimation. *Bioinformatics*, 20(11):1737–1754, 2004.

[94]J. Prados, A. Kalousis, J.-C. Sánchez, L. Allard, O. Carrette, and M. Hilario. Mining mass-spectra for diagnosis and biomarker discovery of cerebral accidents. *Proteomics*, 4(8):2320–2332, 2004.

[95]H. Ressom, R. Varghese, M. Abel-Hamid, S. Abdel-Latif Eissa, D. Saha, L. Goldman, E. Petricoin, T. Conrads, T. Veenstra, C. Loffredo, and R. Goldman. Analysis of mass spectral serum profiles for biomarker selection. *Bioinformatics*, 21(21):4039–4045, 2005.

[96]H. Ressom, R. Varghese, S. Drake, G. Hortin, M. Abel-Hamid, C. Loffredo, and R. Goldman. Peak selection from MALDI-TOF mass spectra using ant colony optimization. *Bioinformatics*, 23(5):619–626, 2007.

[97]D. Ross, U. Scherf, M. Eisen, C. Perou, C. Rees, P. Spellman, V. Iyer, S. Jeffrey, M. Van de Rijn, M. Waltham, A. Pergamenschikov, J. Lee, D. Lashkari, D. Shalon, T. Myers, J. Weinstein, D. Botstein, and P. Brown. Systematic variation in gene expression patterns in human cancer cell lines. *Nature Genetics*, 24(3):227–234, 2000.

[98]R. Ruiz, J. Riquelme, and J. Aguilar-Ruiz. Incremental wrapper-based gene selection from microarray data for cancer classification. *Pattern Recognition*, 39:2383–2392, 2006.

[99]Y. Saeys, S. Degroeve, D. Aeyels, P. Rouzé, and Y. Van de Peer. Feature selection for splice site prediction: a new method using EDA-based feature ranking. *BMC Bioinformatics*, 5(1):64, 2004.

[100]Y. Saeys, P. Rouzé, and Y. Van de Peer. In search of the small ones: improved prediction of short exons in vertebrates, plants, fungi, and protists. *Bioinformatics*, 23(4):414–420, 2007.

[101]S. Salzberg, A. Delcher, S. Kasif, and O. White. Microbial gene identification using interpolated markov models. *Nucleic Acids Research*, 26:544–548, 1998.

[102]S. Scheid and R. Spang. twilight; a Bioconductor package for estimating the local false discovery rate. *Bioinformatics*, 21(12):2921–2922, 2005.

[103]S. Shah and A. Kusiak. Data mining and genetic algorithm based gene/SNP selection. *Artificial Intelligence in Medicine*, 31:183–196, 2004.

[104]P. Sham, S. Ao, J. Kwan, P. Kao, F. Cheung, P. Fong, and M. Ng. Combining functional and linkage disequilibrium information in the selection of tag snps. *Bioinformatics*, 23(1):129–131, 2007.

[105]H. Shin and M. Markey. A machine learning perspective on the development of clinical decision support systems utilizing mass spectra of blood samples. *Journal of Biomedical Informatics*, 39:227–248, 2006.

[106]W. Siedelecky and J. Sklansky. On automatic feature selection. *International Journal of Pattern Recognition*, 2:197–220, 1988.

[107]C. Sima, U. Braga-Neto, and E. Dougherty. Superior feature-set ranking for small samples using bolstered error estimation. *Bioinformatics*, 21(7):1046–1054, 2005.

[108]C. Sima and E. Dougherty. What should be expected from feature selection in small-sample settings. *Bioinformatics*, 22(19):2430–2436, 2006.

[109]S. Sinha. Discriminative motifs. *Journal of Computational Biology*, 10(3-4):599–615, 2003.

[110]D. Skalak. Prototype and feature selection by sampling and random mutation hill climbing algorithms. In *Proceedings of the Eleventh International Conference on Machine Learning*, pages 293–301, 1994.

[111]G. Smyth. Linear models and empirical Bayes methods for assessing differential expression in microarray experiments. *Statistical Applications in Genetics and Molecular Biology*, 3(1):Article 3, 2004.

[112]R. Somorjai, B. Dolenko, and R. Baumgartner. Class prediction and discovery using gene microarray and proteomics mass spectroscopy data: curses, caveats, cautions. *Bioinformatics*, 19(12):1484–1491, 2003.

[113]A. Statnikov, C. Aliferis, I. Tsamardinos, D. Hardin, and S. Levy. A comprensive evaluation of multicategory classification methods for microarray gene expression cancer diagnosis. *Bioinformatics*, 21(5):631–643, 2005.

[114]J. Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society. Series B*, 64:479–498, 2002.

[115]Y. Su, T. Murali, V. Pavlovic, M. Schaffer, and S. Kasif. RankGene: identification of diagnostic genes based on expression data. *Bioinformatics*, 19(12):1587–1579, 2003.

[116]M. Tadesse, M. Vannucci, and P. Lio. Identification of DNA regulatory motifs using Bayesian variable selection. *Bioinformatics*, 20(16):2553–2561, 2004.

[117]J. Thomas, J. Olson, S. Tapscott, and L. Zhao. An efficient and robust statistical modeling approach to discover differentially expressed genes using genomic expression profiles. *Genome Research*, 11:1227–1236, 2001.

[118]R. Tibshirani, T. Hastie, B. Narasimhan, S. Soltys, G. Shi, A. Koong, and Q.-T. Le. Sample classification from protein mass spectrometry, by 'peak probability contrast'. *Bioinformatics*, 20(17):3034–3044, 2004.

[119]V. Trevino and F. Falciani. GALGO: an R package for multivariate variable selection using genetic algorithms. *Bioinformatics*, 22(9):1154–1156, 2006.

[120]O. Troyanskaya, M. Garber, P. Brown, D. Bolstein, and R. Altman. Nonparametric methods for identifying differentially expressed genes in microarray data. *Bioinformatics*, 18(11):1454–1461, 2002.

[121]V. Tusher, R. Tibshirani, and G. Chu. Significance analysis of microarrays applied to ionizing radiation response. In *Proceedings of the National Academy of Sciences*, volume 98, pages 5116–5121, 2001.

[122]R. Varshavsky, A. Gottlieb, M. Linial, and D. Horn. Novel unsupervised feature filtering of biological data. *Bioinformatics*, 22(14):e507–513, 2006.

[123]Y. Wang, F. Makedon, and J. Pearlman. Tumor classification based on DNA copy number aberrations determined using SNPS arrays. *Oncology Reports*, 5:1057–1059, 2006.

[124]Y. Wang, I. Tetko, M. Hall, E. Frank, A. Facius, K. Mayer, and H. Mewes. Gene selection from microarray data for cancer classification - a machine learning approach. *Computational Biology and Chemistry*, 29:37–46, 2005.

[125]J. Weston, A. Elisseeff, B. Schoelkopf, and M. Tipping. Use of the zero-norm with linear models and kernel methods. *Journal of Machine Learning Research*, 3:1439–1461, 2003.

[126]I. Witten and E. Frank. *Data Mining: Practical Machine Learning Tools and Techniques, 2nd Edition*. Morgan Kaufmann, San Francisco, 2005.

[127]B. Wu, T. Abbott, D. Fishman, W. McMurray, G. Mor, K. Stone, D. Ward, K. Williams, and H. Zhao. Comparison of statistical methods for classification of ovarian cancer using mass spectometry data. *Bioinformatics*, 19(13):1636–1643, 2003.

[128]E. P. Xing, M. I. Jordan, and R. M. Karp. Feature selection for high-dimensional genomic microarray data. In *Proceedings of the Eighteenth International Conference on Machine Learning*, pages 601–608, 2001.

[129]M. Xiong, Z. Fang, and J. Zhao. Biomarker identification by feature wrappers. *Genome Research*, 11:1878–1887, 2001.

[130]Y. Yang, Y. Xiao, and M. Segal. Identifying differentially expressed genes from microarray experiments via statistic synthesis. *Bioinformatics*, 21(7):1084–1093, 2005.

[131]E. Yeoh, M. Ross, S. Shurtleff, W. Williams, D. Patel, R. Mahfouz, F. Behm, S. Raimondi, M. Relling, A. Patel, P. A, C. Cheng, D. Campana, D. Wilkins, X. Zhou, J. Li, H. Liu, C. Pui, W. Evans, C. Naeve, L. Wong, and J. Downing. Classification, subtype discovery, and prediction of outcome in pediatric lymphoblastic leukemia by gene expression profiling. *Cancer Cell*, 1:133–143, 2002.

[132]K. Yeung and R. Bumgarner. Multiclass classification of microarray data with repeated measurements: application to cancer. *Genome Biology*, 4(12):R83, 2003.

[133]K. Yeung, R. Bumgarner, and A. Raftery. Bayesian model averaging: development of an improved multi-class, gene selection and classification tool for microarray data. *Bioinformatics*, 21(10):2394–2402, 2005.

[134]J. Yu and X. Chen. Bayesian neural network approaches to ovarian cancer identification from high-resolution mass spectometry data. *Bioinformatics*, 21(Suppl. 1):i487–i494, 2005.

[135]J. Yu, S. Ongarello, R. Fiedler, X. Chen, G. Toffolo, C. Cobelli, and Z. Trajanoski. Ovarian cancer identification based on dimensionality reduction for high-throughput mass spectometry data. *Bioinformatics*, 21(12):2200–2209, 2005.

[136]L. Yu and H. Liu. Efficient feature selection via analysis of relevance and redundancy. *Journal of Machine Learning Research*, 5(Oct):1205–1224, 2004.

[137]N. Zavaljevsky, F. Stevens, and J. Reifman. Support vector machines with selective kernel scaling for protein classification and identification of key amino acid positions. *Bioinformatics*, 18(5):689–696, 2002.

[138]X. Zhang, X. Liu, Q. Shi, X.-Q. Xu, H.-C. Leung, L. Harris, J. Iglehart, A. Miron, J. Liu, and W. Wong. Recursive SVM feature selection and sample classification for mass-spectometry and microarray data. *BMC Bioinformatics*, 7(197), 2006.